
Information Technology - Serial Storage Architecture - Transport Layer 2 (SSA-TL2)

Preliminary draft proposed American National Standard

This is an internal working document of T10, a Technical Committee of Accredited Standards Committee NCITS. As such, this is not a completed standard, but is considered to be technically complete by the T10.1 Task Group. The contents of this document may be modified as a result of comments received during the review process.

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

ABSTRACT

This document describes the transport layer of the Serial Storage Architecture (SSA). SSA defines a serial interface hierarchy to be used for purposes within its distance and performance characteristics, including but not limited to storage subsystems.

The Serial Storage Architecture fills a need in the evolution from parallel to serial interfaces for storage devices. It meets the space constraints and cabling considerations for high-density storage arrays with a commensurate improvement in reliability and configurability. The reliability improvements are derived from an architected error recovery, redundant paths to devices, a wrap mode for self-test, line fault detection and a balanced signaling scheme that achieves a low error rate. The configurability results from the ability to hot-plug devices, the self-configuration capability, the 10 meter length of cable segments. The algorithms for the web topology and transport mechanisms for the protocol services need to be defined in a manner that allows migration of the physical and protocol independently.

Project Leader:
John P. Scheible
IBM Corporation.
Bldg 902 Mailstop 9263
11400 Burnet Road
Austin, TX 78758
Voice: (512) 823-8208
FAX: (512) 838-3822
Email: scheible@vnet.ibm.com

Technical Editor:
John P. Scheible
IBM Corporation.
Bldg 902 Mailstop 9263
11400 Burnet Road
Austin, TX 78758
Voice: (512) 823-8208
FAX: (512) 838-3822
Email: scheible@vnet.ibm.com

Other Points of Contact:

T10.1 Chairman
Lawrence Lamers
Adaptec
MS 293
691 South Milpitas Blvd.
Milpitas CA, 95035

Voice: 408-957-7817
Fax: 408-957-7193
Email: ljlamers@aol.com

T10.1 Vice-Chairman
John P. Scheible
IBM Corporation.
Bldg 902 Mailstop 9263
11400 Burnet Road
Austin, TX 78758
(512) 823-8208
(512) 838-3822
scheible@vnet.ibm.com

NCITS Secretariat
Administrator Standards Processing
NCITS Secretariat
1250 Eye Street, NW Suite 200
Washington, DC 20005

Voice: 202-737-8888
FAX: 202-638-4922
Email: x3sec@itic.nw.dc.us

SSA Reflector

Internet address for subscription to the T10.1 reflector:
Note should contain a line stating...
Internet address for distribution via T10.1 reflector:

majordomo@symbios.com
subscribe X3T10-SSA *your email address*
x3t10-ssa@symbios.com

T10 Bulletin Board

719-533-7950

FTP Site:

[ftp.symbios.com](ftp://ftp.symbios.com/pub/standards/io/x3t10.1)
[/pub/standards/io/x3t10.1](ftp://ftp.symbios.com/pub/standards/io/x3t10.1)

Web sites:

or <http://www.x3.org/x3t10>
<http://www.symbios.com/x3t10>
<http://www.ssaia.org>

Document Distribution
Global Engineering
15 Inverness Way East
Englewood, CO 80112-5704

Voice: 303-792-2181
or: 800-854-7179
FAX: 303-792-2192

PATENT STATEMENT

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard.

As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more claims has been received.

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

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

Revision History

Significant technical changes from SSA-TL2 rev 4 are shown with revision marks.

T10.1/1147D revision 4a (February 16, 1997)

During the February 5, 1997 T10.1 working group, the editor was instructed to generate revision 4a of SSA-TL2 with the following modifications:

- 1) 97a101r3 - Letter ballot comment resolution on SSA-TL2 rev 4

X3T10.1/1147D revision 4 (November 24, 1996)

During the October 30th 1996 X3T10.1 plenary, the editor was instructed to generate revision 4 of SSA-TL2 with the following modifications:

- 1) Add 96a155r0 - Comments on TL2 and S3P (part regarding SAT regions).
- 2) Add 95a168r0 - Corrections to reset sections.
- 3) Add 96a172r2 - Simplifications of Extended Distance Option.

X3T10.1/1147D revision 3 (September 16, 1996)

During the August 28th 1996 X3T10.1 plenary, the editor was instructed to generate revision 3 of SSA-TL2 with the following modifications:

- 1) Add 95a210r3 - Interlocked election of master process.
- 2) Add 96a142r5 - Changes to SSA-PH2 rev 2 proposal (as they affect TL2)
- 3) Add 96a145r3 - TL2 Configurator Table management support.
- 4) Add 96a154r0 - Clarification of resets proposal (as modified).
- 5) Add 96a155r0 - Comments on TL2 and S3P (parts as modified).
- 6) Add 96a161r1 - AER reporting between ULP and TL2 layers.
- 7) Add 96a163r0 - Change format of Configurator Table entry.

X3T10.1/1147D revision 2 (July 23, 1996)

During the June 6th 1996 X3T10.1 plenary, I was instructed to generate revision 2 of SSA-TL2 with the following modifications:

- 1) Add 96a134r0 - Third party data transfers.
- 2) Add 96a143r0 - Resurrection of Longer full speed Link proposal, as modified.
- 3) Add 96a144r0 - Notification of change in ULPs supported.
- 4) Add 96a146r1 - Changes to SSA-TL2 proposal
- 5) Continuing editorial changes.

Although it was the intention of the committee to add a new version of 96a210, that version was not available in time for this version of SSA-TL2.

X3T10.1/1147D revision 1 (May 13, 1996)

During the May 2nd 1996 X3T10.1 plenary, I was instructed to generate revision 1 of SSA-TL2 based off of SSA-TL2 rev 0 with the following modifications:

- 1) Add 95a227r2 - Support for SAT regions.
- 2) Add what to report in the SSA-TL VERSION field when multiple versions supported.
- 3) Define when tag values are active.

X3T10.1/1147D revision 0 (April 15, 1996)

During the February 29th 1996 X3T10.1 plenary, I was instructed to generate revision 0 of SSA-TL2 based off of SSA-TL1 rev 10 with the following modifications:

- 1) Update the boilerplate to SSA-TL2.
- 2) Update the definitions of the CUR20 and CUR40 bits in the QUERY PORT SMS.
- 3) Add 95a175r1 - Buffer threshold in 20/40 MB proposal.
- 4) Add 96a113r3 - Simple speed negotiation.
- 5) Verified conformance to ANSI guidelines (96a114r0).
- 6) I defined the SSA-TL VERSION field as 00h-0Fh reserved for SSA-TL1 and 10h being this standard with 11h-FFh reserved.

American National Standard
for Information Systems -
Serial Storage Architecture -
Transport Layer 2 (SSA-TL2)

Secretariat
Information Technology Industry Council

Approved Month dd, yy

American National Standards Institute, Inc.

Abstract

This document describes the transport layer of the Serial Storage Architecture (SSA). SSA defines a serial interface hierarchy to be used for purposes within its distance and performance characteristics, including but not limited to storage subsystems.

The Serial Storage Architecture fills a need in the evolution from parallel to serial interfaces for storage devices. It meets the space constraints and cabling considerations for high-density storage arrays with a commensurate improvement in reliability and configurability. The reliability improvements are derived from an architected error recovery, redundant paths to devices, a wrap mode for self-test, line fault detection and a balanced signaling scheme that achieves a low error rate. The configurability results from the ability to hot-plug devices, the self-configuration capability, the 10 meter length of cable segments. The algorithms for the web topology and transport mechanisms for the protocol services need to be defined in a manner that allows migration of the physical and protocol independently.

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 judgment 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 in no circumstances gives 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 NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard.

As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more claims has been received.

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

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

Published by

American National Standards Institute

11 W. 42nd Street, New York, New York 10036

Copyright 199n by American National Standards Institute

All rights reserved.

Contents		Page
1	Scope	1
1.1	SSA-TL2 characteristics	1
1.2	SSA family of standards	2
2	Normative references	3
3	Definitions, symbols and abbreviations	3
3.1	Definitions.....	3
3.2	Symbols and abbreviations.....	5
4	Conventions	6
5	Data Rate.....	7
6	Web	7
6.1	Strings	8
6.2	Loops	8
6.3	Complex Web containing switches	9
7	8B/10B code	9
7.1	Data characters	10
7.2	Special characters.....	12
7.3	Initialization.....	14
7.4	Code violations.....	14
7.5	Permitted Special character sequences	14
8	Frames	15
8.1	CONTROL field	16
8.2	STATUS field.....	18
8.3	EXT FSN field	18
8.4	PATH field	18
8.5	ADDRESS field	18
8.6	DATA field	19
8.7	CRC field.....	19
9	Protocol.....	20
9.1	Router functions	20
9.2	Web Routing.....	21
9.3	Node routing	23
9.4	Flow Control	23
9.5	SAT.....	26
9.6	Priority for transmission.....	31
10	Link management.....	31
10.1	Port attributes	31
10.2	Node based data structures.....	38
10.3	Beginning Communication process.....	41
10.4	Speed Negotiation process.....	41
10.5	Configuration.....	42
10.6	Resets.....	43
11	Error handling	46
11.1	Link ERP.....	46
11.2	Invalid SMS field handling.....	52
11.3	Asynchronous Alert process	54
11.4	De-queue Asynchronous Alert process	54
11.5	ASYNC ALERT SMS Handling process	54
11.6	MASTER ALERT SMS Handling process	57
12	SSA message structures.....	58
12.1	SMS validation	59
12.2	Privileged SMS.....	60
12.3	Application SMS.....	81
13	Transport layer processes.....	86
13.1	Master Negotiation process	86

13.2	Master Alive process	89
13.3	SAT Region process	89
13.4	Configurator Table Maintenance process	91

Annexes		Page
A	Transport Layer Data Transfer Processes	93
B	Transport Layer Services	95
C	Examples of half and full duplex	106

Figures		Page
1	Relationship of the SSA standards	2
2	Dual-port node	8
3	A string	8
4	A dedicated connection	8
5	A loop	9
6	A complex Web containing a cyclic path	9
7	Generalized frame format	15
8	Specific frame format	16
9	ADDRESS field	18
10	Path addressing	23
11	Timing of acknowledgments	25
12	SAT operation for a loop	28
13	SAT operation for a string	29
14	Web with 2 SAT regions	29
15	State transitions	33

Tables		Page
1	Bit ordering in a byte	6
2	Sustained data rate and maximum distance	7
3	5B/6B encoding rules	11
4	3B/4B encoding rules	11
5	Special characters	12
6	User Defined character definition	14
7	Frame contents	16
8	CONTROL field	16
9	FRAME TYPE field	16
10	RESET TYPE field	17
11	RR flag values when entering the Enabled state	26
12	RR flag values when entering the Ready state	26
13	RR pending flag control	26
14	Port table entry	36

15	Queued Asynchronous Alert table entry	37
16	Asynchronous Alert Address table entry	37
17	UNIQUE ID field	38
18	Configuration table entry	39
19	Configurator table entry	40
20	Configurator table entry STATUS field	40
21	Link Status Byte	49
22	Receiver errors	49
23	Recovery from corrupted ACK character pair	52
24	Response to non-zero reserved fields or reserved code values	54
25	Generic SSA message structure	58
26	SSA-TL2 messages supported	59
27	ASYNC ALERT SMS	60
28	Format of the ALERT CODE field	60
29	ALERT CODE values	62
30	ASYNC REPLY SMS	63
31	RETURN CODE field values	63
32	CONFIGURE PORT SMS	64
33	MODE field values	65
34	DELETE RETURN PATH ID SMS	66
35	MASTER ALERT SMS	67
36	QUERY NODE SMS	68
37	SSA-TL VERSION field values	68
38	QUERY NODE REPLY SMS	69
39	Upper-level protocol code values	70
40	QUERY PORT SMS	71
41	QUERY PORT REPLY SMS	72
42	SUPPORTED SPEED field	73
43	CURRENT SPEED field	73
44	QUERY PROTOCOL SMS	73
45	QUERY PROTOCOL REPLY SMS	74
46	QUERY REGISTRATION SMS	75
47	Configurator table returned data entry format	75
48	QUERY SWITCH SMS	76
49	QUERY SWITCH REPLY SMS	76
50	QUIESCE SMS	77
51	REQUEST SAT REGION SMS	78
52	SAT OPCODE field	79
53	REQUEST SAT REGION REPLY SMS	80
54	RETURN CODE values	80
55	RESPONSE SMS	80
56	RETURN CODE values	81
57	DATA READY SMS	81
58	DATA REPLY SMS	83
59	DATA REQUEST SMS	84
60	DATA REDIRECT SMS	85

Foreword (This foreword is not part of American National Standard NCITS.nnn-199n.)

This Serial Storage Architecture - Transport Layer 2 standard is designed to provide a transport layer acceptable to vendors, looking for an evolution from parallel SCSI, and systems designers looking for opportunities to more fully exploit the capabilities inherent to a serial bus.

This standard was developed by Task Group T10.1 of Accredited Standards Committee NCITS during 1993-96. The standards approval process started in 1995. The informative annexes contained in this document are not part of the standard.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the NCITS Secretariat, Information Technology Industry Council, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This standard was processed and approved for submittal to ANSI by Accredited Standards Committee on Information Processing Systems, NCITS. Committee approval of the standard does not necessarily imply that all committee members voted for approval. At the time it approved this standard, the NCITS Committee had the following members:

James D. Converse, Chairman

Donald C. Loughry, Vice-Chairman

Joanne M. Flanagan, Secretary

<i>Organization Represented</i>	<i>Name of Representative</i>
American Nuclear Society	Geraldine C. Main Sally Hartzell (Alt.)
AMP, Inc	Edward Kelly Charles Brill (Alt.)
Apple Computer	Karen Higginbottom
Association of the Institute for Certification of Professionals (AICCP)	Kenneth Zemrowski
AT&T/NCR	Thomas W. Kern Thomas F. Frost (Alt.)
Boeing Company	Catherine Howells Andrea Vanosdoll (Alt.)
Bull HN Information Systems, Inc.	William George
Compaq Computer Corporation	James Barnes
Digital Equipment Corporation	Delbert Shoemaker Kevin Lewis (Alt.)
Eastman Kodak	James D. Converse Michael Nier (Alt.)
GUIDE International	Frank Kirshenbaum Harold Kuneke (Alt.)
Hewlett-Packard	Donald C. Loughry
Hitachi America, Ltd.	John Neumann Kei Yamashita (Alt.)
Hughes Aircraft Company	Harold L. Zebrack
IBM Corporation	Joel Urman Mary Anne Lawler (Alt.)
National Communication Systems	Dennis Bodson
National Institute of Standards and Technology	Robert E. Roundtree Michael Hogan (Alt.)
Northern Telecom, Inc	Mel Woinsky Subhash Patel (Alt.)
Neville & Associates	Carlton Neville
Recognition Technology Users Association	Herbert P. Schantz G. Edwin Hale (Alt.)
SHARE, Inc	Gary Ainsworth David Thewlis (Alt.)
Sony Corporation	Michael Deese
Storage Technology Corporation	Joseph S. Zajackowski Samuel D. Cheatham (Alt.)
Sun Microsystems	Scott Jameson

3M Company	Gary Robinson (Alt.) Eddie T. Morioka
Unisys Corporation.....	Paul D. Jahnke (Alt.) John L. Hill
U.S. Department of Defense.....	Stephen P. Oksala (Alt.) William C. Rinehuls
U.S. Department of Energy	C. J. Pasquariello (Alt.) Alton Cox
U.S. General Services Administration	Lawrence A. Wasson (Alt.) Douglas Arai
Wintergreen Information Services	Larry L. Jackson (Alt.) Joun Wheeler
Xerox Corporation	Dwight McBain Roy Peirce (Alt.)

Technical Committee T10 on I/O Interfaces, which reviewed this standard, had the following members:

John B. Lohmeyer, Chairman

Lawrence J. Lamers, Vice-Chairman

Ralph Weber, Secretary

Mr. I. Dal Allan	Mr. Paul D. Aloisi	Mr. David Andreatta
Mr. Tak Asami	Mr. Akram Atallah	Mr. Wayne Baldwin
Mr. Geoffrey Barton	Mr. Robert Bellino	Mr. Rick Bohn
Mr. Paul Boulay	Mr. Charles Brill	Mr. Michael Bryan
Mr. John Cannon	Mr. Ting Li Chan	Mr. Shufan Chan
Mr. Joe Chen	Mr. Andy Chen	Mr. Jack Chen
Ms. Nancy Cheng	Mr. Mike Chennery	Mr. Dan Colegrove
Mr. Roger Cummings	Mr. Chris D'Iorio	Mr. Zane Daggett
Mr. William Dallas	Mr. Joe Dambach	Mr. Brian N. Davis
Mr. Jan V. Dedek	Mr. Dhiru N. Desai	Mr. Mike Eneboe
Mr. Mark Evans	Mr. Timothy Feldman	Mr. Stephen G. Finch
Mr. Edward Fong	Mr. Edward A. Gardner	Mr. John Geldman
Mr. Raymond Gilson	Mr. Chuck Grant	Mr. Louis Grantham
Mr. Dave Guss	Mr. Peter Haas	Mr. Douglas Hagerman
Mr. Kenneth J. Hallam	Dr. William Ham	Mr. Tom Hanan
Mr. Norm Harris	Mr. Edward Haske	Mr. Stephen F. Heil
Mr. Mike Hetzel	Mr. Stephen Holmstead	Mr. Gerald Houlder
Mr. Peter Johansson	Mr. Gerry Johnsen	Mr. Brian Johnson
Mr. Skip Jones	Mr. Richard Kalish	Mr. Greg Kapraun
Mr. Thomas J. Kulesza	Mr. Lawrence J. Lamers	Mr. Dennis Lang
Mr. Edward Lappin	Mr. Pat LaVarre	Mr. Jaff Lin
Mr. Robert Liu	Mr. John Lohmeyer	Mr. Bill Mable
Mr. John Masiewicz	Mr. David McFadden	Mr. James McGrath
Mr. Pete McLean	Mr. Patrick Mercer	Mr. Gene Milligan
Dr. Akira James Miura	Mr. E.J. Mondor	Mr. Charles Monia
Mr. Ian Morrell	Mr. John Moy	Mr. S. Nadershahi
Mr. Jay Neer	Mr. Marc A. Noblitt	Mr. Tim Norman
Mr. Vit Novak	Mr. Erich Oetting	Mr. Dennis Pak
Mr. Duncan Penman	Mr. George Penokie	Mr. Doug Piper
Mr. Kevin R. Pokorney	Mr. Gary Porter	Mr. Robert Reisch
Mr. Ron Roberts	Mr. Frank Samela	Mr. John P. Scheible
Mr. Steve D. Schueler	Mr. J. R. Sims	Mr. Michael Smith
Mr. Robert N. Snively	Mr. Jeff Stai	Mr. Gary R. Stephens
Mr. Arlan P. Stone	Mr. Clifford E. Strang Jr.	Mr. Nicos Syrimis
Mr. Matthew Thomas	Mr. Pete Tobias	Mr. Tokuyuki Totani
Mr. Dennis Van Dalsen	Mr. Joseph Wach	Mr. Dean Wallace
Mr. Roger Wang	Mr. Gary M. Watson	Mr. Ralph O. Weber
Mr. Bob Whiteman	Mr. Jeffrey L. Williams	Mr. Michael Wingard
Mr. Devon Worrell	Mr. Anthony Yang	Mr. Danny Yeung
Mr. Ruben Yomtoubian		

Task Group T10.1 on Serial Storage Architecture, which developed this standard, had the following members:

Lawrence J. Lamers, Chairman

John Scheible, Vice-Chairman

Charles Brill
Wolfgang Drichelt
Bill Ham
John Scheible
Greg Alvey
Adge Hawes
Charles Monia
Herb Silverman

David Deming
Neil Edmunds
Brad Kitson
Kelly Tenuta
Bob Atkinson
Bill Mable
Brian Morissette
Daniel Tsai

Mark DeWilde
Chuck Grant
Lawrence Lamers
Michael Wingard
Robert Bellino
Gary Manchester
Said Rahmani

Introduction

This standard is divided into the following clauses and annexes.

Clause 1 defines the scope of the Serial Storage Architecture - SSA-TL2 (Transport Layer 2).

Clause 2 specifies the normative references.

Clause 3 defines the definitions, symbols and abbreviations.

Clause 4 defines the conventions.

Clause 5 defines the Web.

Clause 6 defines the characters.

Clause 7 defines the frames.

Clause 8 defines the protocol.

Clause 9 defines the link management.

Clause 10 defines the error handling.

Clause 11 defines the SMSs.

Clause 12 defines the processes.

Annex A is informative and defines the data transfer processes.

Annex B is informative and defines the services.

Annex C is informative and shows examples of half and full duplex operation

American National Standard for Information Systems -

Information Technology - Serial Storage Architecture Transport Layer 2 (SSA-TL2)

1 Scope

This document defines a transport layer of the Serial Storage Architecture (SSA) that runs SSA-S2P and SSA-S3P while running on SSA-PH2.

The goals of SSA-TL2 are:

- a) provide an Extended Distance Option;
- b) provide support for higher data rates in the physical layer 2 (SSA-PH2);
- c) enhance packet formats and addressing methods;
- d) define a transport layer acceptable to vendors looking for an evolution from parallel SCSI and systems designers looking for opportunities to more fully exploit the capabilities inherent to a serial bus;
- e) and other capabilities that fit within the scope of SSA-TL2 that may be proposed during the development phase by the participants in the project.

This document defines the transport layer of the Serial Storage Architecture (SSA). SSA defines a serial interface hierarchy to be used for purposes within its distance and performance characteristics, including but not limited to storage subsystems. This standard is intended to be used with an upper layer protocol [e.g., SCSI-2 Protocol (SSA-S2P) or SCSI-3 Protocol (SSA-S3P)] and a physical layer [e.g., SSA Physical Layer 2 (SSA-PH2)].

1.1 SSA-TL2 characteristics

The essential characteristics of SSA-TL2 are as follows:

- a) SSA-TL2 is optimized to minimize gate count while maintaining function and performance.
- b) SSA-TL2 provides a flexible addressing scheme to allow the configuration of a Web into dedicated connections, strings, loops, and complex configurations including a switch.
- c) The Web supports peer-to-peer communication with frame multiplexing between any pair of ports.
- d) The Web is self-configuring without the use of address switches.
- e) The link supports point-to-point physical connections only. Strings and loops are created by using dual port nodes with cut-through routing.
- f) An Extended Distance Option exists which allows 10 km between nodes through the use of additional buffering and an increased ACK time-out.
- g) Allows operation at a distance of 2.5 km at a line data rate of 10 MB/s, and at a distance of 340 m at a line data rate of 40 MB/s. The Extended Distance Option allows operation at a distance of 10 km at a line data rate of 40 MB/s. The actual data rates and distances depend on the cable assembly and the drivers and receivers used.
- h) The link makes optimum use of the physical medium by using full-duplex communication to avoid arbitration overhead and turn-around delay.
- i) Full duplex communication achieves an aggregate 80 MB/sec bandwidth (40 MB/sec in each direction) between two ports.

- j) The frames are constructed with a minimum framing overhead (e.g., only 8 characters of overhead for a 128 byte data field).
- k) An architected error recovery procedure provides transparent frame recovery after a transmission error.
- l) A Wrap mode is provided to allow fault isolation at a node.
- m) An optional line fault detector detects failures in the link.
- n) Cyclical paths allow a Web to be configured providing redundant access to a node.
- o) The transport layer is capable of simultaneously supporting multiple upper-level protocols (e.g., SCSI-2, SCSI-3, IPI-3, TCP/IP, etc.).
- p) The 128-byte frame size minimizes the gate count needed for buffers.

This standard defines the following functions:

- a) the protocol (e.g., framing, addressing and flow control);
- b) link management (e.g., buffering, port states, resets, configuration and error recovery);

1.2 SSA family of standards

The relationships of the SSA family of standards is illustrated in Figure 1.




Figure 1 - Relationship of the SSA standards

1.2.1 SSA-S2P

Serial Storage Architecture - (SSA-S2P) SCSI-2 Protocol defines the SCSI-2 Protocol used with SSA-TL1 or SSA-TL2.

1.2.2 SSA-S3P

Serial Storage Architecture - (SSA-S3P) SCSI-3 Protocol defines the SCSI-3 Protocol used with SSA-TL2.

1.2.3 SSA-TL1

Serial Storage Architecture - (SSA-TL1) Transport Layer 1 defines the Transport layer that supports SSA-S2P and requires SSA-PH1.

1.2.4 SSA-TL2

Serial Storage Architecture - (SSA-TL2) Transport Layer 2 defines the Transport layer that supports SSA-S2P and SSA-S3P and requires SSA-PH1 or SSA-PH2.

1.2.5 SSA-PH1

Serial Storage Architecture - (SSA-PH1) Physical Layer 1 defines the Physical layer that supports SSA-TL1 and SSA-TL2, and consists of the electrical characteristics of the interface and the connectors.

1.2.6 SSA-PH2

Serial Storage Architecture - (SSA-PH2) Physical Layer 2 defines the Physical layer that supports SSA-TL2, and consists of the electrical characteristics of the interface and the connectors.

2 References

2.1 Normative references

The following standards contain provisions which, through reference in SSA-TL2, constitute provisions of this standard. At the time of publication, the edition indicated was valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standard listed below. Members of IEC and ISO maintain registers of currently valid International Standards, and ANSI maintains registers for American National Standards.

ANSI X3.294:1996, *Information Technology, Serial Storage Architecture - SCSI-2 Protocol (SSA-S2P)*

T10.1/1051D, *Information Technology, Serial Storage Architecture - SCSI-3 Protocol (SSA-S3P)*

Editor's note: This clause will be adjusted at time of publication to reference the appropriate standards.

2.2 Informative references

The following related standards are for informational purposes only, and do not contain provisions of this standard. At the time of publication, the edition indicated was valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standard listed below. Members of IEC and ISO maintain registers of currently valid International Standards, and ANSI maintains registers for American National Standards.

X3.293:1996 - *Information Technology - Serial Storage Architecture - Physical Layer 1 (SSA-PH1)*

X3.295:1996 - *Information Technology - Serial Storage Architecture - Transport Layer 1 (SSA-TL1)*

T10.1 / 1146D - *Information Technology, Serial Storage Architecture - Physical Layer 2 (SSA-PH2)*

Editor's note: This clause will be adjusted at time of publication to reference appropriate standards.

3 Definitions, symbols and abbreviations

3.1 Definitions

3.1.1 application: A process that is communicating via the link.

3.1.2 APPLICATION: A possible value of the FRAME TYPE field.

3.1.3 Application frame: A frame with a FRAME TYPE field value of APPLICATION.

3.1.4 Assuming Master: A Pending Master that has received all of the Formal Resignations from the Resigning Masters is the Assuming Master until it has completed the web configuration and is then the web Master. The Assuming Master may receive a Resignation Request and not become the web Master.

3.1.5 channel: The facilities in a port to receive an SMS or a single data transfer.

3.1.6 character: A sequence of 10 encoded bits that represents a data byte or a protocol function.

3.1.7 Configurator node: A node that knows the topology of the Web by performing the Configuration process. If a node is not a Configurator node, then it is a Responder node.

3.1.8 CONTROL: A possible value of the FRAME TYPE field.

- 3.1.9 CONTROL field:** The first byte of a frame.
- 3.1.10 Control frame:** A frame with a FRAME TYPE field value of CONTROL.
- 3.1.11 cut-through frames:** Frames that are not addressed to the receiving node and are passed from the input lines of one port to the output lines of another port with minimal latency according to the routing rules.
- 3.1.12 cut-through routing:** A process where a frame not addressed to a node is routed to the output lines of another port as it streams in the input lines of the receiving port.
- 3.1.13 cut-through router:** A router capable of cut through routing or delayed routing.
- 3.1.14 cyclic path:** A portion of a Web where starting at a given node, it is possible to return to that same node after traversing one or more links without traversing any node more than once. A loop is an example of a cyclic path.
- 3.1.15 Data frame:** An application frame with a non-zero Channel component.
- 3.1.16 delayed routing:** A process where a frame not addressed to a node is buffered and then routed to the output lines of another port at a later time.
- 3.1.17 destination node:** The node where the frame arrives with an path component of 00h.
- 3.1.18 Dual port node:** A node with two ports and a cut through router between them.
- 3.1.19 field:** A group of related data characters in a frame (e.g., the CRC field).
- 3.1.20 Formal Resignation:** A Resigning Master Formally Resigns to a node with which it has completed a Resignation Request transaction at an earlier time when it sends a QUERY NODE SMS with the MA bit set to that node.
- 3.1.21 forward:** The act of receiving a frame not addressed to the node on one port and transmitting out another port after modifying the frame according to the protocol rules.
- 3.1.22 frame:** A sequence of 6 or more data characters surrounded by FLAG characters.
- 3.1.23 line:** A physical connection between a transmitter and a receiver.
- 3.1.24 link:** A serial connection between two ports.
- 3.1.25 loop:** A cyclic Web containing dual port nodes only.
- 3.1.26 Master:** A function performed by a single Configutor unique node within a web that coordinates error recovery, configures the ports in the web, monitors the web for blocked communications and clears them with web resets, and manages SAT regions. A switch may be the web Master if it contains a Configutor function and wins the election of the Master.
- 3.1.27 meaning code:** The description assigned the combination of the TYPE, SUBTYPE, and TYPE INFO fields used as a Link ERP exit, in the ASYNC ALERT SMS or MASTER ALERT SMS.
- 3.1.28 node:** A system, controller, switch or device with one or more ports.
- 3.1.29 Originating frames:** Frames that are generated from the node itself and are sent out the output lines on one of the ports of the node.
- 3.1.30 path:** The links and intermediate nodes used to transfer a frame from the source to the destination.
- 3.1.31 Pending Master:** The current Master of a sub-web that, after walking a newly attached sub web containing one or more active Masters, has discovered that it has won mastership of that new sub web. It waits to assume mastership of the new sub-web until all of the resigning masters formally resign.
- 3.1.32 Pending Resignation:** A Resignation is Pending whenever a QUERY NODE REPLY SMS is sent with the CM bit set. The Pending Resignation is satisfied when the Formal Resignation is performed.
- 3.1.33 port:** The hardware and firmware that implements one end of a link.
- 3.1.34 PRIVILEGED:** A possible value of the FRAME TYPE field.
- 3.1.35 Privileged frame:** A frame with a frame type field value of PRIVILEGED.

- 3.1.36 receiver, Rx:** The logic that decodes the signal on the inbound line.
- 3.1.37 remote port:** The port connected to the other end of the SSA link.
- 3.1.38 RESERVED:** A possible value of the FRAME TYPE field or the RESET TYPE field.
- 3.1.39 reserved field:** A field that is reserved for future standardization.
- 3.1.40 reserved field value:** A field value that is reserved for future standardization.
- 3.1.41 Reserved frame:** A frame with a frame type field value of RESERVED.
- 3.1.42 Reserved Control frame:** A frame with a frame type field value of CONTROL, and a RESET TYPE field value of RESERVED.
- 3.1.43 Resigning Master:** The (temporary) Master node of a Sub-Web that has received one or more Resignation Requests. It no longer accepts ASYNC ALERT SMS s, but has not yet finished processing those received prior to the QUERY NODE SMS.
- 3.1.44 Resignation Request:** Received by a Master or Resigning master of a sub-web, this is a QUERY NODE SMS that contains either a higher master priority than the receiving node s, or the same Master priority and a higher Unique ID. A resignation request is answered with a QUERY NODE REPLY with the CM bit set.
- 3.1.45 Responder node:** A node with no knowledge of the Web topology. When a Responder node receives an SMS requiring a response, the RETURN PATH or RETURN PATH ID is included within the SMS to be used to send the response. If a node is not a Responder node, then it is a Configurator node.
- 3.1.46 source node:** The node that originates a particular frame.
- 3.1.47 SMS:** The DATA field portion of a frame with a channel field of 00h.
- 3.1.48 spatial reuse:** The property where multiple sets of transfers between source and destination pairs are allowed on the Web at any given time.
- 3.1.49 string:** A linear Web of dual port nodes. (The nodes at either end shall be a switch, a single port node or a dual port node with one port not operational.)
- 3.1.50 Sub-Web:** A string of nodes with Responder functions, Configurator functions, or a combination of both that are added to another web as a unit.
- 3.1.51 switch:** A node with 4 or more ports (an even number) capable of routing frames from any port to any other port based on the address field of the frame. A switch node shall contain a Configurator or Responder function in addition to the switch function.
- 3.1.52 transmitter, Tx:** The logic that drives the outbound line.
- 3.1.53 Web:** A collection of SSA nodes addressable by each other that are connected by links (i.e., a loop, a string, or complex configurations).

3.2 Symbols and abbreviations

- CMOS** complementary metal oxide semiconductor.
- CRC** cyclic redundancy check.
- DMA** direct memory access.
- ERP** error recovery procedure.
- FCS** fiber channel standard.
- FDDI** fiber distributed data interface.
- FSN** frame sequence number.
- LED** light emitting diode.
- LSI** large scale integration.
- OUI** Organizationally Unique Identifier
- POR** power-on reset.
- POST** power-on self-test.
- RAS** reliability, availability and serviceability.

- RSC** receive sequence counter.
- SCSI** Small Computer Systems Interface.
- SSA** Serial Storage Architecture.
- SMS** SSA Message Structure.
- TSC** transmit sequence counter.
- &** logical AND
- =** assignment or comparison for EQUAL
- ≠** comparison for NOT EQUAL
- <** comparison for LESS THAN
- ≤** comparison for LESS THAN OR EQUAL TO
- >** comparison for GREATER THAN
- +** ADD
- SUBTRACT
- *** MULTIPLY
- ±** PLUS OR MINUS
- ≈** APPROXIMATELY
- »** MUCH GREATER THAN

4 Conventions

Certain words and terms used in this standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in the glossary or in the text where they first appear. Lower case is used for words having the normal English meaning.

Fields containing only one bit are usually referred to as the "named" bit instead of the "named" field. When a bit is set its value is 1. When a bit is cleared, its value is 0.

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 immediately followed by lower-case "h" (xxh) are hexadecimal values.

Decimal numbers are indicated with a comma(e.g., two and one half is represented as "2,5). Decimal numbers having a value exceeding 999 are represented with a space (e.g., 24 255).

The bit ordering within a byte used in SSA-TL2 is illustrated in Table 1.

Table 1 - Bit ordering in a byte

Bit 7	6	5	4	3	2	1	Bit 0
msb							lsb

The bits in an encoded character are designated *a b c d e i f g h j*. Bit *a* is transmitted on the line first and the other bits follow in the order shown.

Reserved bits, fields, bytes, and code values are set aside for future standardization. Their use and interpretation may be specified by future extensions to this standard. A reserved bit, field, or byte shall be set to zero, or in accordance with a future extension of this standard. Reserved bits, fields, bytes or reserved field values shall be ignored when cut-through routing a frame. A destination node that receives a reserved field value in the FRAME TYPE or RESET TYPE fields of the CONTROL field of a frame shall generate a FRAME REJECT ERROR as defined in 11.1.1.9. A destination node that received an SMS with a reserved bit, field, or byte that is not zero, or receives a reserved code value shall respond as defined in 11.2.7.

Ignored bits, fields, or bytes shall be ignored by the receiving node. Any value shall be considered valid.

SMS names are shown as all capital letters, such as QUERY NODE SMS. Field names are shown as small capital letters, such as the FRAME TYPE field. Field values are shown as all capital letters, such as the FRAME TYPE field APPLICATION FRAME value. Variable names are in italics, such as *N*.

The byte ordering convention is Big Endian (i.e., the most significant byte of a number is sent first).

In case of conflict, figures take precedence over tables, and both figures and tables take precedence over text.

SSA-PH is used to refer to either SSA-PH1 or SSA-PH2. SSA-TL is used to refer to either SSA-TL1 or SSA-TL2.

5 Data Rate

SSA links operate at a data rate of either 40 MB/s (i.e., 400 Mb/s) or 20 MB/s (i.e., 200 Mb/s) on the line.

The minimum overhead imposed by the frame format is 8 characters for a FLAG plus the control, address and CRC fields. If the link is being operated in full-duplex mode there is an additional overhead of 4 characters for the ACK and RR responses. With 128-byte data fields these overheads lead to the sustained data rates tabulated in Table 2. The equations used are as follows.

- a) half duplex = $128/(128+8) * 40 \text{ MB/s}$ (or 20 MB/s)
- b) full duplex = $128/(128+8+4) * 40 \text{ MB/s} * 2 \text{ directions}$ (or 20 MB/s)

Table 2 - Sustained data rate and maximum distance

Nominal speed	Half-duplex	Full-duplex	Maximum distance at full speed	Extended Distance Option
40 MB/s	37,6 MB/s	2 x 36,6 MB/s	340 m	10 km
20 MB/s	18,8 MB/s	2 x 18,3 MB/s	680 m	10 km

These data rates assume that there are no transmission errors and that the transmitters and receivers have sufficient buffering to handle back-to-back frames.

The maximum distances allow the responses to be returned within the period of one frame, assuming a propagation velocity of 2×10^8 meters/second. Longer distances are possible at a reduced data rate. However the minimum $25 \mu\text{s}$ ACK time-out imposes an upper limit of 2.5 kilometers in the standard distance mode. At this distance the maximum sustained data rate is approximately $\frac{1}{4}$ speed. The $100 \mu\text{s}$ ACK time-out of the Extended Distance Option increases the upper limit to 10 km, while increased buffering allows 10 km operate at full speed.

The transmit clock shall be accurate to ± 250 PPM.

The receiver shall acquire bit synchronization from the transitions in the transmitted data. If an asynchronous sampling technique is used then the design shall take metastability into account.

6 Web

An SSA Web consists of two or more nodes interconnected by links. Each node has one or more ports. A port consists of the hardware and firmware to support one end of a link.

A node is identified in several ways.

- a) Each node is either a Configurator or a Responder, and this determines the processes and SMS codes that shall be supported. A Configurator node knows the topology of the Web but a Responder node does not.
- b) At any time, at most one node in the Web is a Master, who is responsible for coordinating error recovery. Every Configurator node shall be capable of being a Master node, but a Responder node shall not be capable of being a Master.
- c) A node has one port (i.e., Single port), two ports (i.e., Dual port), or more than two ports (i.e., Switch).

Dual port nodes and switches contain a router between the ports and the node function. Figure 2 shows a dual port node.




Figure 2 - Dual-port node

Depending on the ADDRESS field, the router forwards an inbound frame to the node function or to the outbound line of the other port. When the dual port node originates a frame it instructs the router to transmit it via the specified port.

The following clauses describe the types of Webs possible.

6.1 Strings

A string is a linear Web of two or more nodes, as shown in Figure 3. The node at either end of a string shall be a single port node, a dual port node with one port not operational or a switch. The intermediate nodes (if any) shall be dual port nodes.




Figure 3 - A string

Each link operates independently, managing its own flow control and error recovery. An error on one link does not affect the other links, except to abort a frame that is already being forwarded.

To allow unrestricted communication between any two nodes the maximum number of nodes in a string is 129, including the end nodes.

The simplest case of a string is a dedicated connection between 2 single port nodes, as shown in Figure 4.




Figure 4 - A dedicated connection

6.2 Loops

A loop is a cyclic Web containing only dual port nodes, as shown in Figure 5. A loop shall not contain a switch because a switch breaks the continuity of the loop.




Figure 5 - A loop

A loop provides better availability than a string because the failure of a dual port node or link does not block the communication between any pair of the remaining nodes. The dynamic insertion or removal of a dual port node does not prevent communication between the other nodes.

To allow the Configuration process to determine that a loop is closed, the maximum number of nodes is limited to 128.

An important property of SSA is spatial reuse. Each link functions independently of the other links so that several transfers occur concurrently at the full bandwidth provided that each uses different links. This allows a much higher through-put than a typical bus.

Fairness algorithms are used to share the link bandwidth equitably when there is contention for the use of a particular link (see 9.5).

6.3 Complex Web containing switches

Switches allow strings to be connected to allow over 65 000 nodes to intercommunicate with full error reporting capability. Switches also allow alternative paths to be provided to achieve fault tolerance. A complex Web including 2 switches and 5 strings, where one string is a cyclic path, is shown in Figure 6.




Figure 6 - A complex Web containing a cyclic path

The maximum number of addressable ports in a switch is 126. However this maximum is reduced when the path preceding the switch is a string. For example, if a string of 40 nodes leads to a switch then the maximum number of switch ports addressable by the first node is 86.

7 8B/10B code

Data bytes and protocol functions are encoded into 10-bit characters for transmission on the physical medium.

Synchronous clocking restricts the bit patterns that the transmitter uses since it is undesirable to have long sequences of consecutive zeros or ones. Therefore an encoding algorithm is required to convert the arbitrary data that one wishes to send into patterns suitable for transmission.

The link uses an 8B/10B code¹. The 8B/10B code provides a (0,4) run-length constraint. This means that the minimum run-length of consecutive zeros or ones is 1 bit and the maximum run-length is 5 bits. The 8B/10B code also provides DC-balance with a maximum digital sum variation of 6.

The 8B/10B mapping is achieved efficiently by partitioning each incoming byte into a 5-bit sub-block and a 3-bit sub-block. The sub-blocks are then encoded with separate 5B/6B and 3B/4B encoders.

In order for the encoded data to be DC-balanced, the combined 8B/10B encoder accumulates a running disparity. As a result, the encoder and decoder require sequential logic in addition to combinatorial logic. The disparity of a sub-block is the difference between the number of 1s and 0s in the encoded data; positive and negative disparities refer to an excess of 1s and 0s respectively. For both the 6B and the 4B sub-blocks the permitted disparity is 0, +2 or -2. The 8B/10B code requires that the polarity (positive or negative) of encoded sub-blocks with non-zero disparity shall alternate. The 4B and 6B encodes with non-zero disparity are allocated in complementary pairs to allow this. With a few exceptions, encoded sub-blocks with zero disparity are generally independent of the running disparity.

7.1 Data characters

Bits 7:3 are encoded into bits a, b, c, d, e, i by the 5B/6B encoder using the rules in Table 3.

The first encode is used when the polarity of the running disparity on entry to the sub-block is as shown in the column D-1. The polarity of the running disparity is equal to the polarity of the most-recent sub-block that had non-zero disparity. An x in column D-1 indicates that the first encode is selected independently of the entry disparity. The column D0 indicates the disparity of the first encode, whose value shall be 0, +2 or -2.

The alternate encode is used if the polarity of the running disparity does not match the entry in column D-1. In this case the encode disparity is the complement of column D0.

¹For informative reference only. A. X. Widmer and P. A/ Franazek, A DC-Balanced, Partition-Block, 8B/10B Transmission Code, IBM Journal of Research and Development, Vol. 27 no. 5, September 1983, pp. 440-450.

Table 3 - 5B/6B encoding rules

Symbol	Input 76543	D-1	Encode abcdei	D0	Alternate abcdei
D.0	00000	+	011000	-	100111
D.1	10000	+	100010	-	011101
D.2	01000	+	010010	-	101101
D.3	11000	x	110001	0	
D.4	00100	+	001010	-	110101
D.5	10100	x	101001	0	
D.6	01100	x	011001	0	
D.7	11100	-	111000	0	000111
D.8	00010	+	000110	-	111001
D.9	10010	x	100101	0	
D.10	01010	x	010101	0	
D.11	11010	x	110100	0	
D.12	00110	x	001101	0	
D.13	10110	x	101100	0	
D.14	01110	x	011100	0	
D.15	11110	+	101000	-	010111
D.16	00001	-	011011	+	100100
D.17	10001	x	100011	0	
D.18	01001	x	010011	0	
D.19	11001	x	110010	0	
D.20	00101	x	001011	0	
D.21	10101	x	101010	0	
D.22	01101	x	011010	0	
D.23	11101	-	111010	+	000101
D.24	00011	+	001100	-	110011
D.25	10011	x	100110	0	
D.26	01011	x	010110	0	
D.27	11011	-	110110	+	001001
D.28	00111	x	001110	0	
D.29	10111	-	101110	+	010001
D.30	01111	-	011110	+	100001
D.31	11111	-	101011	+	010100

Bits 2:0 are encoded into bits f, g, h, j by the 3B/4B encoder using the rules in Table 4.

Table 4 - 3B/4B encoding rules

Symbol	Input 210	D-1	Encode fghj	D0	Alternate fghj
D.0	000	+	0100	-	1011
D.1	100	x	1001	0	
D.2	010	x	0101	0	
D.3	110	-	1100	0	0011
D.4	001	+	0010	-	1101
D.5	101	x	1010	0	
D.6	011	x	0110	0	
D.P7	111	-	1110	+	0001
D.A7	111	-	0111	+	1000

This is interpreted in the same way as the 5B/6B table.

Note that although symbols D.7 in the 5B/6B encodes and D.3 in the 3B/4B encodes have zero disparity they are assigned complementary encodes. This reduces the maximum digital sum variation and eliminates certain undesirable long run lengths.

The encode for bits 2:0 = 111b in the 3B/4B encoder has two possible values. Symbol D.A7 shall replace D.P7 under either of the following conditions:

- a) in the previous sub-block $e = i = 1b$ and the running disparity is negative;
- b) in the previous sub-block $e = i = 0b$ and the running disparity is positive.

7.2 Special characters

The 8B/10B code provides 12 Special characters as defined in Table 5.

Table 5 - Special characters

Symbol	Function	D-1	Encode abcdei fghj	D0	Alternate abcdei fghj
K.28.0 ¹	User defined	-	001111 0100	0	110000 1011
K.28.1 ²	FLAG	-	001111 1001	+	110000 0110
K.28.2	ABORT	-	001111 0101	+	110000 1010
K.28.3	SAT	-	001111 0011	+	110000 1100
K.28.4	SAT'	-	001111 0010	0	110000 1101
K.28.5 ²	DIS	-	001111 1010	+	110000 0101
K.28.6 ¹	User defined	-	001111 0110	+	110000 1001
K.28.7 ²	Invalid	-	001111 1000	0	110000 0111
K.23.7	ACK	-	111010 1000	0	000101 0111
K.27.7	RR	-	110110 1000	0	001001 0111
K.29.7	NUL	-	101110 1000	0	010001 0111
K.30.7 ¹	User defined	-	011110 1000	0	100001 0111
NOTES					
¹ - These symbols are defined in 7.2.9.					
² - The symbol column denotes a comma.					

Three of the special characters contain a sequence (bits a, b, c, d, e, i, f = 001111b) that does not occur in any bit position, either within other characters or in any overlap of characters. They are called Comma characters and their properties make them suitable for establishing byte synchronization.

SSA-TL2 uses 8 of the Special characters as Protocol characters. The special character K28.7 is defined as an Invalid character since K28.7 leads to the detection of a false comma, and shall be handled as a code violation (see 7.4). The other 3 special characters are User Defined characters that are available for use by upper-level protocols.

7.2.1 DIS character

The DIS character is sent when a port is in the Disabled state and is used to establish character synchronization (see 9.1.3.1).

7.2.2 FLAG character

In the Enabled state, the receipt of a FLAG character places the port into the Ready state (see 10.1.3.2). Provided that character synchronization has been achieved the following actions are taken when a FLAG character is detected:

- a) If a data character has been received since the previous FLAG character then the receiver considers the current frame to be complete;
- b) The receiver becomes ready to receive a frame.

FLAG characters are also sent when the link is idle in order to maintain character synchronization by the receiver.

7.2.3 NUL character

The transmitter may insert NUL characters within a frame anywhere after the CONTROL field. A NUL character shall not be sent between the leading FLAG character and the CONTROL field of a frame. Also a NUL character

shall not be sent between frames or in the Disabled state. For example, NUL characters may be inserted within the ADDRESS, DATA and CRC fields of a frame, and within Control frames.

Except for protocol errors, the receiver shall ignore NUL characters by discarding them without changing its state. NUL characters shall not be included when calculating the frame CRC. A node shall insert NUL characters in a cut-through frame when data for that frame is not available from the input port (e.g., NUL characters received).

NUL characters are used in the following cases:

- a) If the transmitter has started to send a frame but the data needed to complete the frame is temporarily unavailable;
- b) If the transmitter is still waiting for an ACK character pair when it is ready to send the trailing FLAG character of the next frame (see 9.4.1).

7.2.4 RR character

The RR characters (Receiver Ready characters) shall always be sent in pairs. An RR character pair is used to pace the transmission of frames, and allows the remote port to send another frame (see 9.4.3).

7.2.5 ACK character

The ACK characters (Acknowledge characters) shall always be sent in pairs. An ACK character pair is used to acknowledge the receipt of a valid frame (see 9.4.1).

7.2.6 ABORT character

A port may abort a frame after it has started to transmit the frame by sending a single ABORT character followed immediately by the trailing FLAG character. The remote port discards the frame as though it had never been sent.

NOTE 1 - If a link error corrupts an ABORT character into a NUL or a FLAG character following an otherwise valid frame, then the receiver accepts the frame as valid. To protect against this it is recommended that the transmitter ensure the accumulated CRC is bad before sending the ABORT character. If the CRC has already been sent then it is recommended that the transmitter send an extra data character before the ABORT character. This extra data character causes the receiving node to use the last three bytes of CRC and the extra byte as the CRC field thereby almost guaranteeing bad CRC is detected. This gives three levels of indicating a bad frame (ABORT characters, bad CRC, and extra byte). However, some hardware implementations leave a small window when bad CRC cannot be guaranteed, so this is an implementers note.

If the ABORT character is received when the frame is already being forwarded, then the ABORT character shall also be forwarded.

7.2.7 SAT character

The SAT character (Satisfied character) is used to balance between spatial reuse and fairness of the Web using the SAT algorithm. The receipt of a SAT character authorizes a port that has not transmitted its quota of frames to hold Cut-through frames and give frames it originates a higher priority as specified by the SAT algorithm (see 9.5).

7.2.8 SAT character

The SAT character is used in string configurations and is immediately forwarded or reflected by the port as part of the SAT algorithm. The SAT character does not influence frame origination (see 9.5).

7.2.9 User Defined characters

SSA-TL2 provides three User Defined characters for use by upper-level protocols. However, to avoid conflicts when more than one upper-level protocol is operating on the same Web, the definition of these characters shall be documented in SSA-TL2 transport layer.

The SYNC character is used as an Index pulse to synchronize multiple devices.

The forwarding of User Defined characters shall be selectively enabled for each port (see 12.2.3).

A User Defined character is processed if the input port is in the Ready or Check states. A User Defined character may be forwarded when the output port is in the Ready or Check states.

A User Defined characters are defined in Table 6.

Table 6 - User Defined character definition

Number	K Code	Definition	Used by
1	K28.0	SYNC	SSA-S2P, SSA-S3P
2	K28.6	reserved	reserved
3	K30.7	reserved	reserved

7.3 Initialization

Until character synchronization has been achieved the running disparity of the transmitter is not known to the receiver. Therefore when the receiver first acquires character synchronization in the Disabled state it shall accept one FLAG or DIS character with either disparity. This character establishes the running disparity in the receiver. To be considered valid subsequent characters shall then follow the disparity rules.

The receiver shall follow a similar procedure to re-initialize its running disparity after it has detected a disparity error.

7.4 Code violations

The receiver shall check for illegal 6B sub-blocks, 4B sub-blocks or the K28.7 characters that are outside of the defined data or Special characters. The receiver shall also check that each sub-block and special character is consistent with the required entry disparity. The error leading to a violation may have occurred in a previous character. Finally, the receiver shall check that the 4B sub-blocks D.P7 and D.A7 are consistent with the previous 6B sub-block.

Code violations are only checked when the port is in the Ready state, and are ignored in all other states. If the receiver detects an error while the node is in the Ready state it shall generate a Link Reset frame with a RECEIVER ERROR field value of CODE VIOLATION in the Link Status Byte (see 11.1.1.5).

7.5 Permitted Special character sequences

This clause defines the restrictions on the transmission of Special characters.

The following pairs of protocol characters are indivisible (i.e., the first character is valid only when followed immediately by the second character):

- a) ACK - ACK,
- b) RR - RR,
- c) ABORT - FLAG.

There shall be at least one data character between an ABORT - FLAG pair and the last preceding FLAG character (e.g., FLAG, ABORT, FLAG is not allowed).

There shall be at least one other character between an RR pair and the last preceding RR pair, and between an ACK pair and the last preceding ACK pair (e.g., RR,RR,RR,RR is not allowed).

Violation of these rules while in the Ready state shall cause the receiver to invoke the Link ERP process and exit with a meaning code of PROTOCOL ERROR (see 11.1.1.6).

The following protocol characters may be inserted anywhere within a frame. They may also be inserted within an idle sequence of FLAG characters except during the first 10 FLAG characters when a port is in the Ready or Check states:

- a) ACK - ACK,
- b) RR - RR,
- c) User Defined characters,

d) SAT and SAT'.

8 Frames

The two ports at opposite ends of a link communicate in units called frames. A frame consists of a sequence of at least 6 data characters delimited by FLAG characters at each end. The FLAG characters that mark the beginning and end of a frame are referred to as the leading and trailing FLAGS respectively. The first data character following a FLAG character marks the beginning of a frame. The last data character preceding a FLAG character marks the end of a frame. A frame is divided into a sequence of 3 or 4 fields as shown in Figure 7.

To minimize overheads the trailing FLAG character may act as the leading FLAG character of the next frame. Therefore to ensure compatibility all nodes shall be able to receive consecutive frames that are separated by a single FLAG.




Figure 7 - Generalized frame format

The maximum lengths of the ADDRESS and DATA fields were chosen to balance between Web size, communication efficiency and implementation cost.

A valid frame is a frame that was not aborted by an Abort character, and does not invoke Link ERP at the receiving port (see 11.1).

The transport layer supports three types of frames: Control, Privileged, and Application frames. The frame type is controlled by the FRAME TYPE field (see 8.1.1).

A Control frame is used for resets. If the length of the DATA field of a Control frame is not zero, the receiving node rejects the frame. A Control frame shall be acted upon immediately by the destination node and is not subject to the pacing rules (see 9.4.3). Control frames are only sent when the transmitting port is in the Ready, Enabled or Check states and in Privileged or Normal mode (see 10.1.3 and 10.1.4).

NOTE 2 - Hardware implementers should allow at least 16 bytes (excluding CRC) of dedicated buffer space for future extensions to the Link Reset and Extended Link Reset frames .

A Privileged frame is used by the transport layer for configuration and error recovery. Privileged frames shall be subject to the pacing rules, and are only sent when the transmitting port is in the Ready state and in Privileged or Normal mode.

An Application frame is used by the transport layer for data transfer and to transport any SMSs sent by the upper-layer protocol. Application frames shall be subject to the pacing rules, and are only sent when the transmitting port is in the Ready state and Normal mode. When a port is in Privileged mode the transmitter and receiver discard Application frames.

The content of a frame depends on the control field as shown in Table 7 and Figure 8.

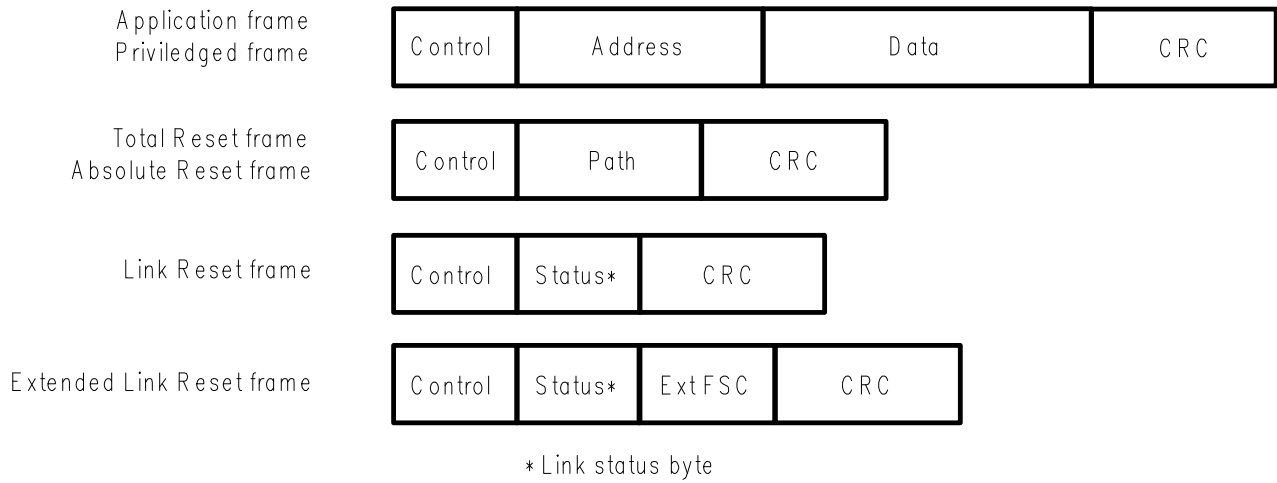


Figure 8 - Specific frame format

Table 7 - Frame contents

CONTROL field		ADDRESS, PATH, STATUS field	DATA field	CRC field
FRAME TYPE	FSN OR RESET TYPE			
APPLICATION	FSN	Path and Channel	Data	CRC
RESERVED				CRC
PRIVILEGED	FSN	Path and Channel	Data	CRC
CONTROL	LINK RESET	Link Status Byte	not present	CRC
CONTROL	EXTENDED LINK RESET	Link Status Byte, Extended frame sequence number	not present	CRC
CONTROL	TOTAL RESET	Path	not present	CRC
CONTROL	RESERVED			CRC
CONTROL	ABSOLUTE RESET	Path	not present	CRC

8.1 CONTROL field

The CONTROL field defined in Table 8 is the first data character following a FLAG character.

Table 8 - CONTROL field

Bit 7	6	5	4	3	2	1	Bit 0
reserved = 0h				FRAME TYPE		RESET TYPE OR FSN	

8.1.1 FRAME TYPE field

The FRAME TYPE field defined in Table 9 specifies the format of the frame.

Table 9 - FRAME TYPE field

Frame type	Description
00b	APPLICATION FRAME
01b	RESERVED
10b	PRIVILEGED FRAME
11b	CONTROL FRAME

8.1.2 RESET TYPE field

The RESET TYPE field specifies the type of reset that shall be performed as a result of receiving a valid Control frame as defined in Table 10.

Table 10 - RESET TYPE field

Reset type	Description
00b	LINK RESET
01b	TOTAL RESET
10b	RESERVED
11b	ABSOLUTE RESET

A Link Reset frame is exchanged between two adjacent nodes to report transmission errors and invoke the Link ERP process (see 10.6.1).

A Total Reset frame causes the destination node to re-initialize its internal state. It is used to recover from a catastrophic error (see 10.6.2).

An Absolute Reset frame causes the destination node to perform a Total Reset and optional additional resets. An Absolute Reset frame optionally invokes the POST (see 10.6.2).

8.1.3 FSN field

The FSN field of a Privileged, Reserved or Application frames contains a FRAME SEQUENCE NUMBER (FSN) that protects against frames being lost by a transmission error. For example, if the sole FLAG character between two frames is corrupted into a valid character then the two frames are merged into one.

The FSN increments modulo 4 in each successive Privileged, Reserved, or Application frame.

Each port shall maintain a 2-bit TRANSMIT SEQUENCE COUNTER and shall copy this into the FSN field of each Privileged, Reserved or Application frame sent. The TRANSMIT SEQUENCE COUNTER is set to 00b in the Disabled state. It is incremented modulo 4 when the trailing FLAG character of a Privileged, Reserved or Application frame is transmitted, regardless of whether the frame is acknowledged. However the TRANSMIT SEQUENCE COUNTER is not incremented if a Control frame is transmitted or the transmitter aborts a frame, as described in 7.2.6.

Each port shall maintain a 2-bit RECEIVE SEQUENCE COUNTER. The RECEIVE SEQUENCE COUNTER is set to 00b in the Disabled state. It is incremented modulo 4 when the trailing FLAG character of a valid Privileged, Reserved or Application frame is received. However the RECEIVE SEQUENCE COUNTER shall not be incremented for a Control frame or an aborted frame.

When a port receives a Privileged, Reserved or Application frame it checks the FSN field against the RECEIVE SEQUENCE COUNTER as follows:

- a) If the FSN field value = RECEIVE SEQUENCE COUNTER, the frame is accepted and an acknowledgment is returned, provided there is no other error.
- b) If the FSN field value \neq RECEIVE SEQUENCE COUNTER then one or more frames have been lost. The current frame shall not be acknowledged and the Link ERP process is invoked resulting in a Link ERP exit with an ALERT CODE value of SEQUENCE ERROR if no other higher priority error exists (see 11.1).

The TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER shall be not incremented for Control frames. The TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER are also used to recover from transmission errors (see 11.1.3). Since the TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER are independent on each link, the router in a dual port or switch node shall regenerate the FSN field as it forwards a frame.

Ports that implement the Extended Distance Option may have to implement additional bits for the TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER, to account for the higher number of unacknowledged frames permitted. A maximum of eight bits per counter are allowed. Only the two least significant bits are reported in the FSN field. The full value of the RECEIVE SEQUENCE COUNTER is exchanged in the EXT RSC field of the Extended Link Reset frame when the Extended Distance Option is enabled.

To facilitate future extensions, dual port and switch nodes shall forward a frame with any bit pattern in the CONTROL field, provided the FSN field value is equal to the RECEIVE SEQUENCE COUNTER for any frames with a FRAME TYPE field value of APPLICATION FRAME, RESERVED or PRIVILEGED FRAME. However the destination port shall invoke the Link ERP process if the FRAME TYPE field is RESERVED or the RESET TYPE field is RESERVED resulting in a Link ERP process exit with an ALERT CODE value of FRAME REJECT ERROR for any reserved bit patterns if no other higher priority error exists (see 11.1).

8.2 STATUS field

The STATUS field begins with the next data character following the CONTROL field in a Link Reset frame and contains the Link Status Byte (see 11.1.2).

8.3 EXT RSC field

The EXT RSC field is a one byte field immediately follows the STATUS field in an Extended Link Reset frame, and contains the extended receive sequence counter. The Extended Link Reset frame is only used for link recovery when the Extended Distance Option is enabled for a link.

8.4 PATH field

The PATH field begins with the next data character following the CONTROL field in Total Reset or Absolute Reset frames. The PATH field is of variable length and consists of a Path component (see 8.5.3).

8.5 ADDRESS field

The ADDRESS field defined in Figure 9 begins with the next data character following the CONTROL field of a Privileged, Reserved or Application frame. The ADDRESS field is of variable length and consists of two components, the Path component and the Channel component (see 8.5.3 and 8.5.4).




Figure 9 - ADDRESS field

The ADDRESS field in a Privileged, Reserved or Application frame contains from 2 to 6 bytes, depending on the complexity of the Web and the number of Channels implemented by the destination node. It consists of a Path component followed by a Channel component. Each component shall be a minimum of 1 byte and is extended a byte at a time by setting the EXTEND bit.

8.5.1 EXTEND bit

When the EXTEND bit is set in a byte it indicates that the component (Path or Channel) is extended into the following byte. When the EXTEND bit is cleared it indicates that the current byte is the last byte of the component. Thus, within the ADDRESS field, the first byte with the EXTEND bit cleared specifies the end of the Path component and the second byte with the EXTEND bit cleared specifies the end of the Channel component.

8.5.2 INDEX

The INDEX is a 7-bit unsigned integer used by the Path and Channel components.

8.5.3 Path component

The Path component is used to route the frame over a particular path to the destination node. Path addresses are relative to the source node. This simplifies the routing hardware and avoids the need to assign absolute node addresses at power-on.

When a frame is forwarded by the router in a dual port node the first byte of the Path component is decremented. When a frame is forwarded by a switch the first byte of the Path component shall be decremented or deleted according to the rules specified in 9.2.3. If it is deleted the next byte shall be used for subsequent routing decisions.

The length of the Path component is 1 - 4 bytes. The maximum length is not checked explicitly. However, it needs to be specified for other functions (e.g., frame buffers and SMS formats).

8.5.4 Channel component

The Channel component selects a channel within the receiving port at the destination node. A channel consists of the facilities to receive an SMS or to receive a single data transfer. Every port shall provide channel zero to receive SMSs. All other Channels are dynamically allocated.

The transport layer starts and controls data transfers by exchanging SMS frames between the source node and the destination node (see 12.3). The destination node allocates a Channel to receive the Data frames and it indicates the number of bytes that it expects.

The length of the Channel component is 1-2 bytes. The maximum length is not checked explicitly. However, it needs to be specified for other functions (e.g., frame buffers and SMS formats). A single-byte Channel component allows up to 127 data Channels; 2 bytes allow an additional 16,256 data Channels.

8.6 DATA field

The DATA field shall be present in all frames except Control frames. The contents of the DATA field in Privileged SSA-TL2 SMSs is defined in 12.2. The contents of the DATA field in SSA-TL2 Application SMSs is defined in 12.3.

The DATA field starts in the next data character following the ADDRESS field. The destination node locates the start of the DATA field by parsing the ADDRESS field [i.e., the remaining Path byte and the Channel byte(s)].

The DATA field contains from 0 to 128 data characters. The maximum length is defined to ensure compatibility between different nodes in a Web.

Except for the Data frame containing the last byte of a data transfer, the length of the DATA field of any Data frame shall be an integer multiple of 8 bytes. This restriction on the division of a data transfer into frames eases alignment at the parallel interface to the node function. A frame that violates the restriction is optionally rejected by the destination node.

A frame shall contain at most one SMS. The maximum length of the DATA field in an SMS is 32 bytes.

8.7 CRC field

The CRC field consists of the 4 data characters that immediately precede the trailing FLAG character. It is accumulated over the CONTROL, PATH, ADDRESS, LINK STATUS BYTE and DATA fields. The router in dual port and switch nodes may interpret the CONTROL, PATH, and ADDRESS fields as they are received (see 9.1). However, the receiving port shall not assume that any of the fields are valid until the CRC field has been determined to be valid.

The value of the CRC field for a frame shall not be calculated to include any interleaved Special characters.

Each node shall generate the CRC field when sending a frame. Each node shall regenerate the CRC field when forwarding a frame.

8.7.1 Polynomial definitions

The generator polynomial is $G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$.

NOTE 3 - SSA, FDDI and FC-PH use the same degree 32 polynomial.

The degree 31 polynomial with all coefficients equal to 1 is $L(x) = x^{31} + x^{30} + x^{29} + x^{28} + \dots + x^3 + x^2 + x + 1$.

The remainder polynomial produced by the receiver in the absence of errors is $C(x) = x^{32} L(x)$ modulo $G(x)$
 $= x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^9 + x^6 + x^5 + x^4 + x^3 + x + 1$.

The degree $k-1$ polynomial representing the CONTROL, ADDRESS and DATA fields of the transmitted frame is $F(x)$.

The remainder polynomial is of degree less than 32 is $R(x)$.

The polynomial of degree less than 32 representing the transmitted CRC field is $K(x)$.

The transmitted frame including the CRC field is represented by $M(x)$.

All arithmetic is Modulo 2.

8.7.2 CRC generation

The CRC field is generated with $R(x) = [x^{32} F(x) + x^k L(x)]$ modulo $G(x)$. This means that the CRC generator shall be preset to all ones at the start of each frame.

The transmitted CRC field is the ones complement of $R(x)$ with $K(x) = L(x) + R(x)$.

The CRC field appended to the transmitted frame is defined by $M(x) = x^{32} F(x) + K(x)$.

8.7.3 Checking

The received frame is divided by $G(x)$. However, to allow for leading zeros the CRC register is first preset to all ones. In the absence of errors the remainder is $C(x) = x^{32} [M(x) + x^k L(x)]$ modulo $G(x)$. If the CRC remainder is not C704DD7Bh the port shall invoke the Link ERP process with a RECEIVER ERROR field value of CRC ERROR (see 11.1.1.7).

9 Protocol

9.1 Router functions

This clause describes in more detail the functions of the router in a dual port or switch node.

Each node in the path between the source node and destination node forwards the frames with minimum delay achieved by using a cut-through router, rather than store-and-forward. Cut-through routing allows each node to begin forwarding a frame out another port while still receiving characters of the same frame, provided that the out-bound port is not already sending a frame or the WAITING FOR RR FLAG is incremented.

The minimum latency introduced by a cut-through router is 5 characters since the CRC field is regenerated as a frame is forwarded. The inbound port shall receive the trailing FLAG character before the first CRC character is sent on the outbound port.

NOTE 4 - In a practical implementation a delay of up to 12 characters is possible. For a string of 16 nodes this allows a minimum delay of less than 10 μ s to be achieved at 20 MB/s. With store-and-forward routing the minimum delay is over 100 μ s with 128-byte frames.

Dual-port nodes shall use a cut-through router when both ports are running at the same speed to ensure good performance in all applications. Dual-port nodes with both ports running at different speeds should delay cutting-through from the slower port to the faster port until enough of the frame is received to allow the faster port to send it without using NUL characters. Switch nodes may also use a cut-through router.

If a node detects an error in an inbound frame when it has already started to forward the frame, it shall send an ABORT character immediately followed by a FLAG character. The ABORT character causes the destination node to discard the preceding frame without generating a link error. If the destination node has also started to forward the frame then it terminates the forwarding frame with an ABORT character immediately followed by a FLAG character.

9.2 Web Routing

Each type of node uses different rules to interpret the Path component in the ADDRESS or PATH fields of an inbound frame. If the receiving port determines that the Path component is invalid according to its rules then it shall invoke the Link ERP process with a RECEIVER ERROR value of FRAME REJECT (see 11.1.1.9).

The following nomenclature is used to define the addressing rules:

- a) First byte refers to the first byte in the Path or Channel component.
- b) Second byte refers to the second byte in the Channel component.
- c) Input port refers to the port of a switch node that receives the frame.
- d) Output port refers to the selected output port of a switch node.
- e) Ports refers to the total number of ports on a switch node.

9.2.1 Single port node

A single port node has only one port, and shall interpret the Path component of the ADDRESS or PATH fields as follows:

```
If First byte = 00h then Do;
    Accept the frame;
    Interpret the remainder of the ADDRESS field as a Channel component;
    End;
Else reject the frame;
```

If the first byte of the Path component is non-zero the node shall invoke the Link ERP process with a RECEIVER ERRORS value of FRAME REJECT (see 11.1.1.9).

For error reporting and configuration, the PORT field has a value of one (see 12.2.1 and 12.2.3).

9.2.2 Dual port node

A dual port node has 2 ports and an internal router. Depending on the value of the first byte of the Path component of the ADDRESS or PATH fields, the router either forwards the frame on the out-bound line of the other port, or accepts the inbound frame for processing internally. The router interprets the Path component as follows:

```
If First byte = 00h then Do;
    Accept the frame;
    Interpret the remainder of the ADDRESS field as a Channel component;
    End;
Else If First byte.Index > 0 then Do;
    First byte.Index = First byte.Index - 1;
    Forward the frame via the other port;
    End;
Else reject the frame;
```

If the first byte of the Path component is 80h the node shall invoke the Link ERP process with a RECEIVER ERRORS value of FRAME REJECT (see 11.1.1.9).

For error reporting and configuration the PORT field has a value of one or two (see 12.2.1 and 12.2.3).

9.2.3 Switch node

A switch node shall have an internal router and an even number of ports in the range of 4 to 126 inclusive. If the first byte of the Path component is zero, the switch shall accept the inbound frame from any port for processing internally. The frame is routed out the selected output port. The switch shall use one of two different addressing algorithms, depending on the relationship of the first Path byte to the number of switch ports.

The first addressing scheme shall be identified by the first Path byte index having a value less than or equal to the number of switch ports. An inbound frame shall be forwarded to the output port identified by the first Path byte index. The first Path byte shall then be deleted. In this scheme, a first Path byte with the EXTEND bit cleared is

a frame reject case, since the last Path byte shall be deleted and the frame forwarded without a Path. This addressing scheme allows for a more flexible port selection, but since the first Path byte is deleted, switches shall only be nested three deep.

The second addressing scheme shall be identified by the first Path byte index having a value greater than the number of switch ports. For routing purposes the switch ports are grouped into pairs. Port 1 is paired with port 2, port 3 is paired with port 4, etc. In general port $2N-1$ is paired with port $2N$ for $N = 1, 2, 3, \dots 63$. An inbound frame shall be forwarded through its paired output port. The first Path byte shall then be decremented by the number of switch ports plus one without deleting it. This provides an efficient means to cascade a number of switches.

The switch router interprets the path component of the address field as follows:

```

If First byte = 00h then do;          (** Process the frame internally **)
    Accept the frame;
    Interpret the remainder of the ADDRESS field as a Channel component;
end;
else if (First byte.Index > Ports)
then do;          (** Decrement and route out paired port **)
    First byte.Index = First byte.Index - (Ports + 1);
    Forward the frame via other port of pair;
end;
else if [(Input port = First byte.Index) or (First byte.Extend = 0b)
or (First byte = 80h)]
    then Reject the frame
else do;          (** Delete and route out selected port **)
    Output port = First byte.Index;
    Delete First byte;
    Forward the frame via Output port;
end;

```

The Link ERP process with a RECEIVER ERRORS value of FRAME REJECT (see 11.1.1.9) shall be invoked for the following conditions:

- a) First byte = 80h (i.e., the frame is addressed to the switch, but another Path component byte exists);
- b) First Byte.index \leq Ports and Extend = 0b (i.e., the first byte shall be discarded, but no additional bytes exist);
- c) or First Byte.index \leq Ports and First byte.Index = Receiving port (i.e., an attempt has been made to output the frame via the receiving port).

Switch ports shall be numbered sequentially starting at one (see 12.2.1 and 12.2.3).

9.2.4 Routing example

An example of a complex Web is shown in Figure 10. The numbers in the edges of the boxes are the port numbers. For illustration purposes the ports are numbered in a clockwise direction. The number inside each box is the hexadecimal Path address of that node relative to the indicated Configurator node. The dual port nodes on the cyclic path allow two alternative paths from the Configurator node. When a particular path includes a switch it has a shorter alias address if the switch input and output ports are in the same pair.

All Configurator nodes execute the Configuration process (see 10.5.1) to discover the ports that are Operational. A Configurator node also selects a primary Path component to each other node and calculates the corresponding Return Path.




Figure 10 - Path addressing

9.3 Node routing

The Channel component is interpreted by all nodes as shown below.

If First byte.Index = 0 then Do;

If First byte.Extend = 0b then select SMS Channel;

Else reject the frame;

End;

Else Do;

If First byte.Extend = 0b then

select Data Channel = First byte.Index;

Else If Second byte.Extend = 0b then

select Data Channel = (128 * First byte.Index) + Second byte.Index;

Else reject the frame;

End;

Channels 0 - 127 are addressed by one byte. Channels 128 - 16383 are addressed by two bytes.

If the first byte of the Channel component is 80h or the Channel component is in error, the node shall invoke the Link ERP process with a RECEIVER ERRORS value of FRAME REJECT (see 11.1.1.9).

9.4 Flow Control

To implement the necessary flow control, the remote port sends the local port character pairs for each frame it receives:

- a) Acknowledgment is a pair of consecutive ACK characters;
- b) Receiver Ready is a pair of consecutive RR characters.

The Special characters are used in pairs to protect the Acknowledgment or Receiver Ready from being manufactured by transmission errors. A node shall only act on an Acknowledgment or Receiver Ready pair when it has received both characters of the pair without any other intervening characters.

Flow control operates on each link independently. ACK and RR characters shall not be forwarded by a router.

In full-duplex operation a port may send an Acknowledgment for a received frame or a Receiver Ready if buffer space is available, while it is transmitting another frame. In this case the transmitter gives priority to the

Acknowledgment or Receiver Ready and interleaves it within the frame. This scheme reduces latency and the amount of buffering required by each port to sustain the full data rate of the link.

9.4.1 Acknowledgments

The link protocol requires a port to acknowledge every valid received frame, except for Total Reset, Reserved or Absolute Reset Control frames. A frame is valid if it does not contain any of the errors listed in 11.1.1. The destination port transmits an ACK character pair when it receives a valid frame. When the source port receives the ACK character pair, it shall discard the transmitted frame from its frame buffers.

This protocol guarantees that the transmitter always associates each ACK character pair unambiguously with the corresponding frame independently of propagation delays, the transmission speed and the frame length. When the Extended Distance Option is active, a transmitter cannot unambiguously associate an ACK character pair with the corresponding frame, and therefore shall retransmit all unacknowledged frames during error recovery.

Each port has two associated flags, WAITING FOR ACK FLAG and ACK PENDING FLAG, that control acknowledgments. Ports that are not capable of the Extended Distance Option shall implement these flags as single bits. Ports that are capable of enabling the Extended Distance Option shall implement the WAITING FOR ACK FLAG and ACK PENDING FLAGS as counters, with enough bits (up to 8) to count the number of transmit and receive frame buffers, respectively, the port provides. The flags operate as follows:

- a) When a port enters the Disabled state it sets the WAITING FOR ACK FLAG and ACK PENDING FLAG to zero.
- b) A port shall increment the WAITING FOR ACK flag after it finishes transmitting the trailing FLAG character of any frame except Total Reset, Reserved or Absolute Reset Control frames. A port decrements the WAITING FOR ACK FLAG when it receives an ACK character pair. The local port shall discard the corresponding transmit data. The error conditions shall be handled as follows:
 - 1) After incrementing the WAITING FOR ACK FLAG a port shall wait to receive the ACK character pair for a minimum of 25 μ s and a maximum of 50 μ s before setting the ACK bit in the Link Status Byte and invoking the Link ERP process (this range allows a simple binary counter to be used for the timer) for the default SSA link behavior. Ports capable of running the Extended Distance Option shall use a 100 μ s to 200 μ s timer, and reset this timer each time an ACK character pair is received while the WAITING FOR ACK flag is greater than zero. Extended distance nodes come out of reset with their ACK time-outs set between 100 μ s and 200 μ s. If an extended distance capable port is connected to a port not capable of extended distance, then the Master shall issue a CONFIGURE PORT SMS with the RACK bit set causing the extended distance capable node to reduce its ACK time-out to a minimum of 25 μ s and a maximum of 50 μ s. In order to guarantee that a port does not indicate an ACK TIME-OUT error for a frame, the delay from the port incrementing the WAITING FOR ACK FLAG to receiving the second ACK character shall be less than 25 μ s for the normal operating mode, and 100 μ s for the Extended Distance Option. This limits the maximum operating distance of the link (see 5)

If a non-Extended Distance Option enabled port's WAITING FOR ACK FLAG is equal to 1 when it finishes transmitting the CRC field of the next frame, then it shall not transmit the trailing FLAG character. Instead it sends NUL characters until either the ACK character pair is received or an ACK time-out occurs. If an ACK time-out occurs in this state, then the port shall abort the frame (see 7.2.6), set the ACK bit in the Link Status Byte and invoke the Link ERP process. Ports that are enabled for the Extended Distance Option may transmit the trailing FLAG character of the next frame even when there are ACKs outstanding. The WAITING FOR ACK FLAG shall keep count of the frames that remain unacknowledged.
 - 2) If a port receives an ACK character pair when the WAITING FOR ACK flag is zero, or it receives only a single ACK character, then it shall set the PROTOCOL ERROR value in the RECEIVER ERRORS field of the Link Status Byte and invokes the Link ERP process.
- c) A port increments the ACK PENDING FLAG when it receives the trailing FLAG character of a valid frame in either of the following cases:
 - 1) the port is in the Ready state and the frame is not a Total Reset, Reserved or Absolute Reset frame;
 - 2) the port is in the Check state and the frame is a Link Reset frame.

- d) When the ACK PENDING FLAG is non-zero the port shall transmit an ACK character pair as soon as possible, but not immediately following the previous ACK character pair. The ACK PENDING flag is decremented when the ACK character pair has been transmitted.

9.4.2 Acknowledgment example

The timing requirements for acknowledgments and the use of NUL characters for non-Extended Distance Option operation is illustrated in Figure 11, where port 1 and port 2 are on opposite ends of a link (i.e., on different nodes). Port 1 finishes sending a first frame and then it immediately sends a second short frame. Due to the delay before port 2 acknowledges the first frame port 1 has to insert a NUL character in the second frame.

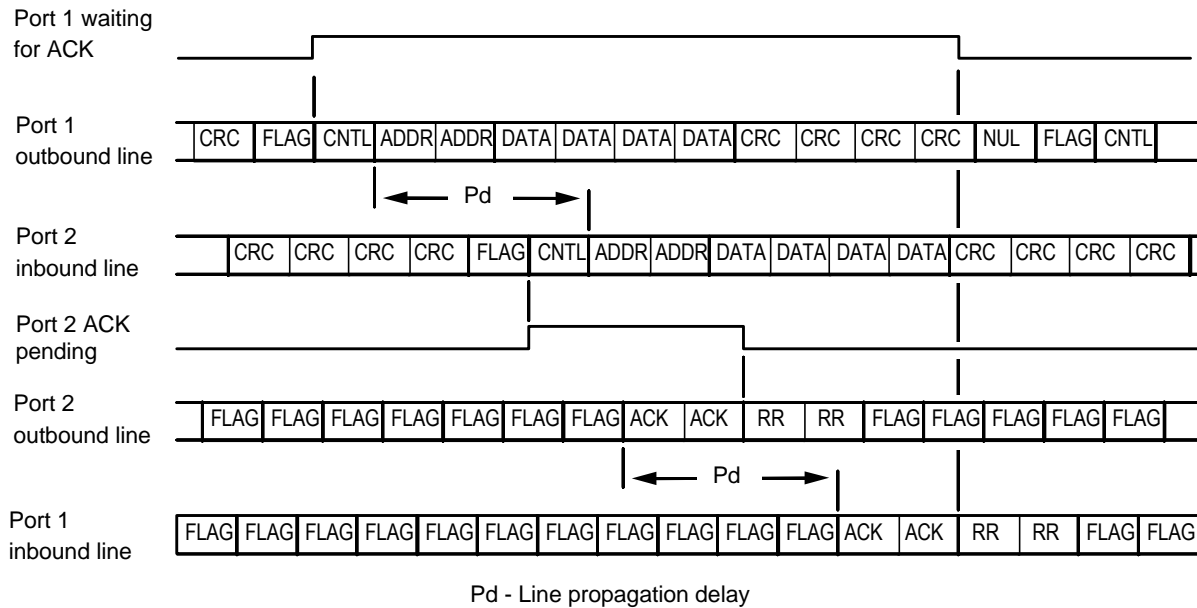


Figure 11 - Timing of acknowledgments (non-Extended Distance Option operation)

9.4.3 Receiver Ready Pacing

Pacing ensures that a transmitting port does not overrun the available buffer space in the receiving port. The unit of pacing is a frame. The receiver shall have buffer space for at least 1 frame, although additional buffer space is required to sustain the full data rate of the link.

Each port has two associated flags, the WAITING FOR RR FLAG and the RR PENDING FLAG, that control pacing. Ports that are not capable of the Extended Distance Option shall implement these flags as single bits. Ports that are capable of enabling the Extended Distance Option shall implement the WAITING FOR RR FLAG as an 8-bit counter, and the RR PENDING FLAG as a counter with enough bits (up to 8) to count the number of receive frame buffers the port provides, and additionally a sign bit to allow the counter to go negative. The flags operate as follows:

- a) When a port enters the Enabled state it sets the WAITING FOR RR FLAG and RR PENDING FLAG according to the Table 11, depending on whether the port is capable of Extended Distance operation, and whether or not the Extended Distance operation is enabled;
- b) A port may start sending a frame when either of the following conditions is satisfied:
 - 1) the port is in the Ready state and its WAITING FOR RR FLAG is set as shown in Table 12;
 - 2) the transmitted frame is a Control frame and the port is in the Enabled, Check or Ready states.
 - 3) the immediately preceding character it transmitted was not an RR.

- c) The WAITING FOR RR FLAG is incremented when a port transmits the CONTROL field of a Privileged, Reserved, or Application frame and is decremented when a port received an RR character pair.
- d) When all of the following conditions are satisfied, a port shall transmit an RR character pair after the current character:
 - 1) the port is in the Ready state and the RR PENDING FLAG is greater than zero;
 - 2) buffer space is available to receive at least one frame on any active Channel, or shall be available before an incoming frame is received;
 - 3) the immediately preceding character transmitted was not an RR character.
- e) The RR PENDING FLAG controlled as defined in Table 13, depending on whether the port is operating with the Extended Distance Option enabled or not.

When a port capable of Extended Distance operation has the option enabled, it shall increment the RR PENDING FLAG by the number of extra receive buffers it has, i.e. it shall consider those buffers to have become available for use.

If an Extended Distance operating port subsequently has the option disabled by a CONFIGURE PORT SMS, then it shall consider the extra receiver buffers to be removed from service, and decrement the RR PENDING FLAG by that amount. If this causes the RR PENDING FLAG to become negative, then the port shall not transmit any more RR character pairs until the RR PENDING FLAG is again greater than zero.

Table 11 - RR flag values when entering the Enabled state

	For non Extended Distance capable port	Extended Distance capable but not enabled	Extended Distance capable and enabled
WAITING FOR RR FLAG	1b	FFh	FFh
RR PENDING FLAG	1b	1	Number of receiver frame buffers available

Table 12 - RR flag values when entering the Ready state

	For non Extended Distance capable port	For Extended Distance capable port
WAITING FOR RR FLAG	0b	less than FFh;

Table 13 - RR pending flag control

RR PENDING FLAG	For non Extended Distance operation	For Extended Distance operation
Incremented when	Port receives the CONTROL field of a Privileged, Reserved or Application frame	Receive buffer becomes available
Decrementd when	Port transmits an RR character pair	Port transmits an RR character pair, or receive buffer is removed from service

9.5 SAT

The origination of Application frames for strings and loops is controlled by the SAT algorithm² where SAT refers to Satisfied. The SAT algorithm helps balance between spatial reuse of the individual links and any node being starved of bandwidth.

9.5.1 Terminology

Some SAT terms are defined in the following list:

²For Informative reference only. I. Cidon and Y. Ofek, MetaRing: A Full-Duplex Ring with Fairness and Spatial Reuse , IEEE INFOCOM 90, pp. 969-981.

- a) A QUOTA field in the CONFIGURE PORT SMS (see 12.2.3) is used to determine how many frames a port may originate while holding cut through frames.
- b) B QUOTA field in the CONFIGURE PORT SMS is used to determine the number of frames a port may originate if no cut through traffic exists (i.e., the link is idle with two consecutive FLAG characters).
- c) REFLECT bit in the CONFIGURE PORT SMS is one of the factors to instruct the port to reflect SAT and SAT characters, breaking the loop or string into two SAT regions (see 9.5.5).
- d) REFLECTION FLAG is a node function that determines whether all ports of a node reflect or forward the SAT and SAT characters (see 9.5.2).
- e) *Hold quota* is a variable containing the A QUOTA field value times the number of current outstanding data transfers from the port. This dynamic scaling of the A QUOTA ensures fairness when a port is originating several outbound data transfers concurrently. The number of frames sent does not have to be balanced between each of the concurrent data transfers.
- f) *Idle quota* is a variable containing the B QUOTA field value times the number of current outstanding data transfers from the port. This dynamic scaling of the B QUOTA ensures fairness when a port is originating several outbound data transfers concurrently. The number of frames sent does not have to be balanced between each of the concurrent data transfers.
- g) FRAME COUNTER is a counter that counts the number of Application frames originated since the last time the SAT character was forwarded or reflected by the port. Forwarding or reflecting a SAT character sets the FRAME COUNTER to zero.
- h) Originated means that the trailing flag of an originating frame has been sent. Aborted frames are not considered Originated and thus not counted against the SAT quotas.
- i) Satisfied refers to the condition of a port whereby either of the two following conditions is met:
 - 1) The port has originated at least *Hold quota* number of frames since it previously forwarded or reflected a SAT character;
 - 2) The port has no frames to originate when holding the SAT character.

9.5.2 Reflection versus forwarding

If the REFLECTION FLAG is cleared, all port(s) of a node shall forward SAT and SAT characters. If the REFLECTION FLAG is set, all port(s) of a node shall reflect SAT and SAT characters.

A port shall forward a SAT character after receiving a SAT character by holding it according to the rules of the SAT algorithm and then generating a SAT character on the outbound lines of the port that did not receive the SAT character. A dual port node shall unconditionally forward a SAT character after receiving a SAT character by generating a SAT character on the outbound lines of the port that did not receive the SAT character.

A port shall reflect a SAT character after receiving a SAT character by holding it according to the rules of the SAT algorithm and then generating a SAT character on the outbound lines of the same port that received the SAT character. A SAT character is reflected unconditionally by generating a SAT character on the outbound lines of the same port that received the SAT character.

9.5.3 The SAT algorithm

The node at the end of a string shall be a single port node, a dual port node with one port not operational or a switch (a dual port node with both ports operational is not at the end of a string by definition). Thus all nodes shall implement the SAT algorithm.

The SAT algorithm is described initially for a loop of dual port nodes, with enhancements for a string shown in 9.5.4. Application frames entering a loop are controlled by a SAT character that circulates in the opposite direction to the frames that it controls, as shown in Figure 12. Circulating the SAT character in the opposite direction maximizes spatial reuse.

The SAT algorithm is applied independently to clockwise and counter-clockwise traffic. Therefore there are two SAT characters circulating on the loop, one in each direction.




Figure 12 - SAT operation for a loop

When a port receives a SAT character it forwards or reflects the character via the other port to the next node if it is Satisfied. Otherwise the port holds the character until it is Satisfied, even if it is waiting for RR. Then it forwards or reflects the SAT character and sets the FRAME COUNTER to zero.

Privileged and Control frames are not controlled by the SAT algorithm and, subject to the link protocol, they may be sent at any time. A port may originate an Application frame, subject to the link protocol, under either of the following conditions:

- a) The port is holding the SAT character and the port is not Satisfied;
- b) There is no frame waiting to be forwarded from another port, and *Idle quota* is greater than FRAME COUNTER.

The ability of a port to originate frames is broken into two cases (a and b), and one case (c) not able to originate frames:

- a) $\text{FRAME COUNTER} < \text{Hold quota}$.
Before the SAT character arrives, the port shall originate frames only when there are no cut through frames pending. After the SAT character arrives, the port shall hold cut through frames until Satisfied, then the SAT character is forwarded or reflected and the FRAME COUNTER shall be set to zero;
- b) $\text{FRAME COUNTER} \geq \text{Hold quota}$ and $\text{FRAME COUNTER} < \text{Idle quota}$.
Before the SAT character arrives, the port shall originate frames only when there are no cut through frames pending. When the SAT character is received, it shall be forwarded or reflected and the FRAME COUNTER shall be set to zero;
- c) $\text{FRAME COUNTER} \geq \text{Idle quota}$.
The port shall not originate any frames. When the SAT character is received, it shall be forwarded or reflected, the FRAME COUNTER shall be set to zero and Application frames may be originated according to the SAT algorithm.

Correct operation of the SAT algorithm requires $\text{B QUOTA} \geq \text{A QUOTA}$, otherwise a node is never Satisfied and it does not forward or reflect the SAT character.

A port shall be able to receive SAT and SAT' characters while it is in any state, provided that character synchronization is maintained. However the port shall only transmit SAT or SAT' characters when in the Ready state.

9.5.4 Extensions of the SAT algorithm for strings

The extensions to the algorithm for a string are as follows:

- a) If a port at the end of a string receives a SAT character then it reflects the SAT character as a SAT' character traveling in the opposite direction after it is Satisfied (see Figure 13).
- b) A dual port node with both ports operational forwards SAT' received from either port unconditionally.

If a port at the end of a string receives a SAT' character then it reflects the SAT' character as a SAT character traveling in the opposite direction immediately without resetting the FRAME COUNTER




Figure 13 - SAT operation for a string

The reflection of SAT as SAT' characters and SAT' as SAT characters is dynamic. If a string of dual port nodes is connected into a closed loop then SAT' characters are no longer generated. To avoid a previous SAT' character circulating indefinitely a dual port node with both ports operational shall discard SAT' characters as follows:

- a) a SAT' character received on port 1 is discarded if no SAT character has been received on port 2 since the previous SAT' character was received on port 1;
- b) a SAT' character received on port 2 is discarded if no SAT character has been received on port 1 since the previous SAT' character was received on port 2.

9.5.5 SAT regions

In some applications where each Configutor node talks exclusively with a subset of the Responders, the traffic on a string or loop is naturally divided into separate regions, as illustrated in Figure 14. In this example, significant traffic is only carried between the following nodes: Configutor1 - Responder1, Configutor1 - Responder2, Configutor2 - Responder3 and Configutor2 - Responder4.

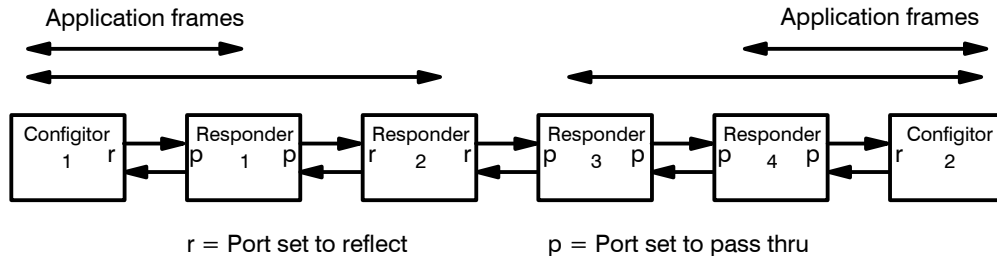


Figure 14 - Web with 2 SAT regions

In this application, any trade-off between fairness and spatial reuse is avoided by logically dividing the Web into 2 regions, each with its own SAT algorithm. This is achieved by setting the REFLECT bit in the CONFIGURE PORT SMSs for either port on Responder 2 (or Responder 3). The ports on Responder 2 then reflects a SAT character as a SAT' character and a SAT' character as a SAT character. Responder 2 still behaves as a dual port node for frame routing (e.g., to allow each system to check the operation of the other).

This facility allows 100% spatial reuse between regions and perfect fairness in each region if A QUOTA=B QUOTA.

If one Configutor node fails the remaining system takes over all of the Responders after issuing a CONFIGURE PORT SMS with new REFLECT, A QUOTA, and B QUOTA fields (if needed).

9.5.6 Timing requirements of the SAT algorithm

All ports incorporate a SAT TIME-OUT timer to detect loss of the SAT character due to a link error. The SAT TIME-OUT timer is set to zero when the port forwards a SAT character or reflects a SAT character as a SAT character. If the SAT TIME-OUT timer counts 100 ms, then a SAT character is created and treated as if it has been received. Multiple SAT characters circulating in the same direction results if several ports generate new SAT characters simultaneously. However the fairness algorithm still operates correctly and the characters eventually merge into a single character as described below.

If a port receives a SAT character while it is holding the previous SAT then it discards the second SAT character.

Each of the following functions shall not introduce a delay exceeding 10 character periods per node:

- a) Forwarding a SAT character when the port is Satisfied, subject to the delay described in the paragraph below;
- b) Reflecting a SAT character as a SAT' character when the port is Satisfied;
- c) Forwarding a SAT' character;
- d) Reflecting a SAT' character as a SAT character.

For a loop or string with only a few nodes the SAT character rotates very rapidly if all nodes are Satisfied. This impacts frames being transmitted in the other direction. Therefore ports shall perform the following two delays:

- a) A port shall not forward a SAT character until at least 100 character periods have elapsed since it forwarded the previous SAT character;
- b) A port shall not reflect a SAT character as a SAT' character until at least 100 character periods have elapsed since it reflected the previous SAT character.

9.5.7 Configuring the quotas

The A QUOTA and B QUOTA values for each port are allocated by the Master node in a CONFIGURE PORT SMS. The optimum values depend on the total number of nodes in each string or loop and the traffic pattern. This clause describes how the values are determined.

The SAT ROTATION time is the time taken by the SAT character to circulate once around a string or loop, as measured in character periods. SAT ROTATION time is a minimum when all ports are Satisfied and they forward the SAT characters immediately (SAT' characters are always forwarded immediately). SAT ROTATION time increases when the Web becomes heavily loaded and ports hold the SAT character.

The SAT DELAY time is the minimum time to propagate or reflect SAT or SAT' characters through a single node whose inbound port is Satisfied, as measured in character periods.

For a string with N nodes the minimum SAT ROTATION time is $2*N*SAT\ DELAY$ characters. The corresponding value for a loop is only $N*SAT\ DELAY$ characters since the propagation delay for SAT' character is eliminated. SAT ROTATION time is always greater than 100 character periods. However after a link failure the loop becomes a string and the minimum SAT ROTATION time doubles. Therefore it is advisable to always allow for a minimum SAT ROTATION time of $2*N*SAT\ DELAY$ time.

The performance advantages of SSA are due in part to its ability to frame-multiplex SMSs and data with low latency. For optimum frame multiplexing and spatial reuse the SAT ROTATION time is minimized. This applies both under light load (when most nodes are Satisfied) and under heavy load (when nodes hold the SAT character).

For long strings SAT DELAY time is a major contributor to SAT ROTATION. Consequently SAT DELAY time has been specified as 10 characters maximum. The SAT and SAT' characters have a high priority for transmission (see 9.6).

The FRAME LENGTH is the total length of a frame. This is typically 136 characters for a data frame. To avoid throttling when a single node wants to originate data frames, $B\ QUOTA \geq (SAT\ ROTATION\ time / FRAME\ LENGTH) \geq (2*N*SAT\ DELAY\ time / FRAME\ LENGTH)$ For the extreme case of a string with 129 nodes this requires a minimum B QUOTA of 19 frames.

In general there is a compromise between fairness and spatial reuse. Setting $B \text{ QUOTA} = A \text{ QUOTA}$ guarantees perfect fairness but it limits spatial reuse. Choosing $B \text{ QUOTA} \gg A \text{ QUOTA}$ permits maximum spatial reuse and therefore maximum total through-put. However in some scenarios this favors the active node that is furthest upstream since cut through traffic normally has priority. Therefore the upstream node is able to originate $B \text{ QUOTA}$ frames for each rotation of the SAT character, whereas each down-stream node only originates $A \text{ QUOTA}$ frames. A reasonable compromise for most applications is $B \text{ QUOTA} \approx 4 * A \text{ QUOTA}$.

9.6 Priority for transmission

When a port has multiple characters or frames eligible for transmission, the priority order shall be as follows:

- a) (Highest priority) A FLAG character if 10 consecutive FLAG characters have not yet been transmitted after entering the Ready state;
- b) Originate the second character of a pair of protocol characters (i.e., ACK-ACK, RR-RR, or ABORT-FLAG);
- c) Transmit any pending SAT or SAT' character(s);
- d) Transmit any pending User Defined character(s);
- e) Originate an ACK character pair;
- f) Originate an RR character pair;
- g) Continue transmitting the current frame (Including NUL and ABORT characters);
- h) Forward a Total Reset, Reserved or Absolute Reset Control frame;
- i) Originate a Control frame.

Transmit an Application or Privileged frame. Dual port nodes normally give priority to frames being forwarded from the other port. This reduces the latency for long strings and loops. However when the transmitter in a dual port node is holding the SAT token it gives priority to frames being originated by the node itself. This improves fairness by reducing the SAT rotation time (see 9.5.7).

NOTE 5 - It is recommended that switch nodes ensure equal fairness between frames being forwarded by the node itself, the other port of a pair and by other ports.

- j) (Lowest priority) Transmit FLAG or DIS characters when idle.

Note that each item in this list has associated pre-requisite conditions before it is eligible for transmission (e.g., Application and Privileged frames cannot be transmitted if the WAITING FOR RR FLAG is equal to 1).

10 Link management

10.1 Port attributes

The following clauses describe those attributes associated with a port.

10.1.1 Frame buffers

Since the unit of pacing is a frame each port shall provide buffering to receive at least 1 frame. The buffer shall accommodate the maximum-length CONTROL, ADDRESS and DATA fields (i.e., $1+6+128 = 135$ bytes total).

In certain occurrences of a transmission error the Link ERP process retransmits the last frames sent. In nodes implementing the Extended Distance Option, this may be up to 255 frames (i.e., the limit of the 8 bit Extended FSC passed during Link ERP). Therefore a port shall also retain each frame it transmits until it receives the corresponding ACK character pair.

To achieve continuous communication at the full bandwidth of the link it is necessary for each port to have additional buffering. If a port has a pair of transmit buffers available and a pair of receive buffers to provide A/B buffering, one buffer of each pair is emptied/filled by the link while the other is filled/emptied by the application. If the Extended Distance Option is implemented, then the additional buffers permit the sending of additional frames while waiting for earlier transmitted frames to be acknowledged. The number of buffers (B) required to maintain full speed data transmission depends on the data rate (R_{DATA}), the round trip propagation delay of the transmission medium (T_{RT}), and the size of the frame (S_{FRAME}). The relationship is as follows:

$$B \geq (T_{RT} * R_{DATA} / S_{FRAME}) + 1$$

NOTE 6 - To maintain full speed (40 MB/s) at full size data frames (128 bytes of data) at the maximum distance (10 KM) require at least 30 buffers under the assumptions of clause 5. Additional dedicated buffers are optional for SMSs to avoid impacting data transfers due to processing delays.

The number and type of buffers in a port are vendor specific. However the following requirements shall be met to ensure performance and compatibility with other nodes in a Web:

- a) All ports shall be able to sustain data transfers with 128 byte DATA fields at the full bandwidth of the link;
- b) When there is no contention, the Router in dual port or switch nodes shall be able to forward the maximum length frames continuously;
- c) All ports shall have an additional buffer to hold a Control frame when all other buffers are full.

10.1.2 TRANSMIT POINTER and RETRY POINTER

The transmit buffering shall be managed to allow correct recovery after a link error. In the case of recovered errors, the Link ERP process retransmits or discards the last one or two frames that were transmitted before an error. Each port shall maintain sufficient status to identify the buffers containing these frames and the order that they were transmitted.

For example, if there are N transmit buffers and the transmitter always accesses them in a cyclic sequence then the following two pointers provide sufficient information:

- a) The TRANSMIT POINTER points to the buffer that is to be transmitted next. It is incremented by 1 modulo N each time the transmitter sends a trailing FLAG character without a preceding ABORT character;
- b) The RETRY POINTER points to the next buffer to be acknowledged. It is incremented by 1 modulo N each time an ACK character pair is received while the WAITING FOR ACK FLAG is non-zero. Normally the RETRY POINTER follows the TRANSMIT POINTER closely but when an error occurs it is possible to lag by 1, or the number of buffers implemented (Extended Distance Option).

The above scheme is an example used later in the description of the Link ERP process. However, sometimes it is impossible to guarantee cyclic use (e.g., if some buffers are reserved for SMSs). In this case, an alternate solution is to keep a log of the buffers that have been transmitted but not acknowledged, and the appropriate actions to be taken when TRANSMIT POINTER and RETRY POINTER are referenced.

10.1.3 Port states

Each port shall be in one of four states as shown in Figure 15. The current state is inspected by the node processor. State changes are controlled by the node processor or change automatically when certain events occur.

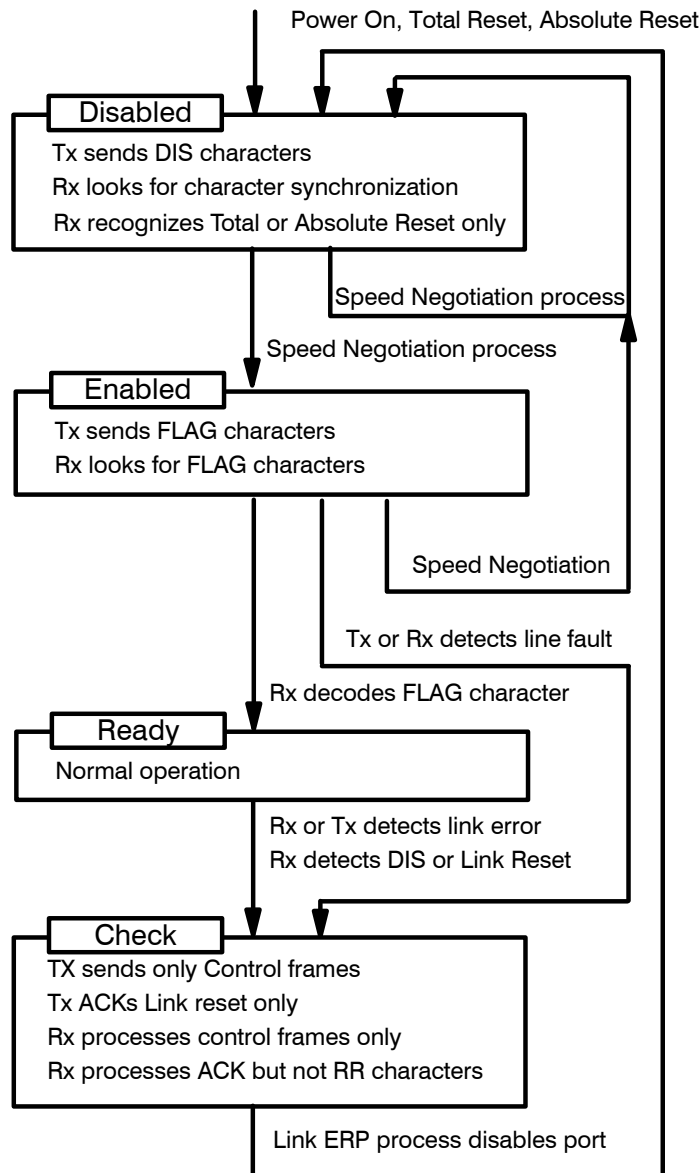


Figure 15 - State transitions

10.1.3.1 Disabled state

The Disabled state is entered after power on, after processing an Absolute Reset or Total Reset frame before the link is made Operational, or is entered during the Link ERP process. Upon entering the Disabled state, the following port functions are initialized:

- a) the LINK RESET FLAG is cleared;
- b) the TRANSMIT SEQUENCE COUNTER is set to 00b;
- c) the RECEIVE TRANSMIT NUMBER is set to 00b;
- d) the RR PENDING FLAG and WAITING FOR RR FLAG are set to one;
- e) the ACK PENDING FLAG and WAITING FOR ACK FLAG are zero;
- f) and all of the error latches are cleared.

The port functions that are not initialized include the following:

- a) the Frame buffers;
- b) the TRANSMIT POINTER;

- c) the RETRY POINTER;
- d) the port mode;
- e) the Port table;
- f) the Queued Asynchronous Alert table;
- g) and the Asynchronous Alert Address table.

In the Disabled state the transmitter sends DIS characters, maintaining DC balance and allows the remote receiver to maintain character synchronization during the Link ERP process. The receiver attempts to acquire character synchronization. If synchronization is achieved the node responds only to a Total Reset, Reserved or Absolute Reset Control frame.

The minimum duration of the Disabled state is 200 character periods. This allows sufficient time to acquire character synchronization when beginning communication since the remote port cannot enter the Ready state until the local port changes to the Enabled state and sends FLAG characters. It also guarantees that the remote node recognizes the Disabled state during the Link ERP process.

The LINK RESET RECEIVED FLAG shall be cleared in the Disabled state (see 10.1.9).

10.1.3.2 Enabled state

The Enabled state is a transient state that is entered from the Disabled state prior to entering the Ready state.

When a port is ready to begin communications the node processor changes the port to the Enabled state (see 9.5). In the Enabled state the transmitter shall transmit FLAG characters and the receiver shall look for FLAG characters. When the receiver decodes a FLAG character the port shall enter the Ready state.

10.1.3.3 Ready state

The Ready state is the state where normal transmission and reception of frames occurs.

In order to allow the remote port sufficient time to enter the Ready state, after a port enters the Ready state it shall transmit at least 10 FLAG characters before sending any other character.

10.1.3.4 Check state

The Check state is entered automatically from the Ready state when a port detects an error or it receives a Link Reset frame. It shall also be entered from the Enabled state when the line driver/receiver detects a line fault. The transition to the Check state invokes the Link ERP process. Except for Control frames, the link is inoperable until the port returns to the Ready state.

When a port enters the Check state the transmitter stops sending Privileged, Reserved, and Application frames after completing or aborting the current frame, if any. The transmitter then sends FLAG characters continuously, except in the following cases:

- a) acknowledge a Link Reset frame;
- b) send a Link Reset frame;
- c) send a User Defined character;
- d) or Abort a Link Reset frame.

When a port is in the Check state the receiver shall discard any incoming Privileged, Reserved or Application frames without generating an ACK or RR character pair. The receiver shall discard RR character pairs but ACK character pairs shall be accepted and acted upon.

10.1.4 Port modes

Independent of the state (see 10.1.3) a port operates in one of three modes, as described in the following clauses.

10.1.4.1 Normal mode

Subject to the port state, Normal mode allows any type of frame to be sent and received.

Normal mode is entered from Privileged mode when the port processes a CONFIGURE PORT SMS with a MODE field value of SET NORMAL MODE to the port (see 12.2.3). The Master node shall be capable of selecting Normal mode internally.

10.1.4.2 Privileged mode

In Privileged mode, a port may transmit and receive Control and Privileged frames and shall discard Application and Reserved frames. When the transmitter discards a frame it shall not increment the TRANSMIT SEQUENCE COUNTER. The receiver discards Application and Reserved frames. When the receiver discards a frame, the port shall generate ACK and RR character pairs normally and shall increment the RECEIVE SEQUENCE COUNTER.

Privileged mode is entered from Wrap mode on exit from the power-on self-test (POST) and from Normal mode when the Link ERP process exits with an ALERT CODE value. Privileged mode is entered from any mode after a Total Reset, Reserved or Absolute Reset Control frame is addressed to the node. Privileged mode is entered when the Master sends a CONFIGURE PORT SMS with a MODE field value of SET PRIVILEGED MODE. The Master shall be capable of selecting Privileged mode internally.

10.1.4.3 Wrap mode

If the port is in Wrap mode, the transmitter output of a port shall be internally connected to the receiver input of the same port. In Wrap mode, frames and their responses share the same line. Except for the line driver and receiver, Wrap mode allows a port to be fully tested without needing a remote node. This is useful to test the local hardware during the POST.

When the port is in Wrap mode, the outbound line shall be held at logic zero and any signal on the inbound line shall be ignored. The output of the line fault detector is ignored.

The port shall be placed in Wrap mode at power-on. On exit from the POST, the port enters Privileged mode, Disabled state and the Beginning Communications process is started. The port is placed in the Wrap mode when the Master sends a CONFIGURE PORT SMS with a MODE field value of SET WRAP MODE. The Master shall be capable of entering and exiting the Wrap mode internally.

When the local port enters the Wrap mode, the remote port at the other end of the link invokes the Link ERP process. The Link ERP process exits unsuccessfully and the remote port optionally invokes the Asynchronous Alert process, depending on how the remote port was configured.

When a port is in the Wrap mode, the OPERATIONAL FLAG shall be cleared. If a received frame is addressed to the port it is discarded and the receiving port invokes an Asynchronous Alert process.

Care shall be taken when placing a port in Wrap mode anytime after the POST since in some configurations if the node processor hangs it is impossible to reset it without a power off-on sequence.

NOTE 7 - Optionally, a node detects when it is not receiving characters on any of its ports and activates a Service indicator (typically an LED). The Master activates the Service indicator on a failing node by wrapping all ports on the remote nodes that are connected to it (the Master node issues CONFIGURE PORT SMSs to the remote nodes or it wraps its own ports internally).

10.1.5 OPERATIONAL FLAG

The OPERATIONAL FLAG is set if the port is able to queue a Privileged frame for transmission to the remote node.

The OPERATIONAL FLAG is cleared after a power-on reset. A port sets the OPERATIONAL FLAG when it enters the Ready state and is not in Wrap mode. When a port changes the OPERATIONAL FLAG from cleared to set, it shall invoke the Asynchronous Alert process with an ALERT CODE value of PORT NOW OPERATIONAL.

The OPERATIONAL FLAG is cleared when the Link ERP process exits, the port is placed in the Wrap mode or the port is reset. The OPERATIONAL FLAG remains set while the port is performing the Link ERP process or when the Link ERP process completes successfully.

The configuration process uses the contents of the QUERY NODE REPLY SMS to determine whether each port of the addressed node has the OPERATIONAL FLAG set or cleared.

A node shall not originate any frame other than a Total Reset, Reserved or an Absolute Reset Control frame via a port that has the OPERATIONAL FLAG cleared.

The reception of a frame, (whether the frame is Acknowledged or not) does not depend on the receiving port s OPERATIONAL FLAG . For example, an Application frame is received when the port is in the Ready state, in Normal mode and the RR PENDING FLAG is zero.

The router in a dual port or switch node forwards frames as follows:

- a) An Application or Privileged frame is only forwarded if the addressed port s OPERATIONAL FLAG is set. Otherwise the node invokes the Asynchronous Alert process with an ALERT CODE value of ADDRESSED PORT NOT OPERATIONAL (see 12.2.1), and the router discards the frame. The receiving port acknowledges the frame and remains in the Ready state.
- b) A Total Reset, Reserved or Absolute Reset Control frame is forwarded regardless of the state of the addressed port s OPERATIONAL FLAG.

10.1.6 TRANSMIT SEQUENCE COUNTER and RECEIVE SEQUENCE COUNTER

The port shall maintain two 2-bit counters, the TRANSMIT SEQUENCE COUNTER and the RECEIVE SEQUENCE COUNTER, to detect lost frames (see 8.1.3).

10.1.7 WAITING FOR RR FLAG and RR PENDING FLAG

The port shall maintain two flags, WAITING FOR RR FLAG and RR PENDING FLAG, to control the pacing of frames (see 9.4.3). If the Extended Distance Option is not supported, then these flags may be single bits. If the port implements the Extended Distance Option, the WAITING FOR RR FLAG shall be 8 bits, and the RR PENDING FLAG shall contain enough bits to count the number of receive buffers implemented, along with a sign bit.

10.1.8 WAITING FOR ACK FLAG and ACK PENDING FLAG

The port maintains two flags, WAITING FOR ACK FLAG and ACK PENDING FLAG, to control the acknowledgment of frames (see 9.4.1). If the Extended Distance Option is not supported, then these flags may be single bits. If the port implements the Extended Distance Option, the WAITING FOR ACK FLAG shall contain enough bits to count the number of transmit buffers implemented. The ACK PENDING FLAG shall contain enough bits to count the number of transmit buffers implemented.

10.1.9 LINK RESET RECEIVED FLAG

The LINK RESET RECEIVED FLAG shall be used during the Link ERP process to determine if the process was entered due to a link error (cleared) or due to receiving a Link Reset frame (set). The LINK RESET RECEIVED FLAG shall be cleared when the port enters the Disabled state. The LINK RESET RECEIVED FLAG shall be set when the port receives a Link Reset frame and is not in the Disabled state. The LINK RESET RECEIVED FLAG is tested during the Link ERP process.

10.1.10 Port table

The node is responsible for maintaining the information necessary to Invoke the Asynchronous Alert process. The information is supplied by the CONFIGURE PORT SMS on a per port basis.

The format of the Port table is vendor specific. However the Port table shall contain the information defined in Table 14 either directly or indirectly. There is one entry for each port.

Table 14 - Port table entry

Field	Description
AA VALID FLAG	A flag indicating if the port table entry is valid
AA PORT	The port to be used for the ASYNC ALERT SMS
AA RETURN PATH	The path to be used for the ASYNC ALERT SMS frame
AA TAG	The TAG field value to be used for the ASYNC ALERT SMS

When the node powers up, the Port table shall be initialized with an entry for each port and with the AA VALID FLAG cleared. When a CONFIGURE PORT SMS is received, the Port table entry associated with the PORT field specified in the CONFIGURE PORT SMS is set as follows:

- a) The AA VALID FLAG is set;
- b) The port number the CONFIGURE PORT SMS was received on is loaded into AA PORT field;
- c) The RETURN PATH field from the CONFIGURE PORT SMS is loaded into AA RETURN PATH field;
- d) The TAG field from the CONFIGURE PORT SMS is loaded into AA TAG field.

10.1.11 Queued Asynchronous Alert table

The node is responsible for maintaining the information necessary for queuing Asynchronous Alerts.

The format of the Queued Asynchronous Alert table is vendor specific. However, the Queued Asynchronous Alert table shall contain the information defined in Table 15 either directly or indirectly. Each entry contains a queued Asynchronous Alert.

Table 15 - Queued Asynchronous Alert table entry

Field	Description
AA ALERT CODE	A three byte value consisting of the ASYNC ALERT SMS ALERT CODE value
AA CONTROL	A 1 byte field containing the CONTROL field of the frame responsible for the Asynchronous Alert. It is zero if not applicable.
AA CHANNEL	A 2 byte field containing the CHANNEL field of the frame responsible for the Asynchronous Alert. It is zero if not applicable.
AA FRAME	A 21 bytes field containing the first 21 bytes of the associated frame s DATA field, left justified and zero filled.

A separate Queued Asynchronous Alert table is maintained for each port, containing an entry for each Asynchronous Alert to be reported against that port. The port to be used to report the Asynchronous Alert is found in the Port table (see 10.1.10).

10.1.12 Asynchronous Alert Address table

The Master is responsible for maintaining the information necessary for associating ASYNC ALERT SMSs with a particular node.

The format of the Asynchronous Alert Address table is vendor specific. However, the Asynchronous Alert Address table shall contain the information defined in Table 16 either directly or indirectly. A table entry exists for each tag value given for an ASYNC ALERT SMS. When the Master issues a CONFIGURE PORT SMS, an entry is created or updated in the Master Asynchronous Alert table. When an ASYNC ALERT SMS is received, the table is used to identify the node and port affected by the Asynchronous Alert. Another possible implementation places this information into the Configuration Table, but the table defined below allows quick access to the information, especially if the TAG value of the ASYNC ALERT SMS is the entry number.

Table 16 - Asynchronous Alert Address table entry

Field	Description
AA TAG	A two byte value consisting of the Asynchronous Alert TAG value
AA ADDRESS	A four byte value consisting of value to be used for the Path component of the ADDRESS field of any ASYNC REPLY SMS sent to the node.
AA PORT	A one byte value indicating the port to use to send the ASYNC REPLY SMS.
AA NODE IDENTIFIER	A pointer into the Configuration table to associate the Asynchronous Alert with a particular node.

10.1.13 SAT TIME-OUT TIMER

The SAT TIME-OUT TIMER shall be implemented by all ports and used to determine that a SAT character has been lost. The SAT TIME-OUT TIMER is used in the SAT algorithm (see 9.5.3) and is a 100 ms timer.

The node shall activate the SAT TIME-OUT TIMER when the port completes Power On, Total Reset, or Absolute Reset. The port shall restart the SAT TIME-OUT TIMER when the port forwards or reflects a SAT character. When the SAT TIME-OUT TIMER expires, the port shall create a SAT character and treats it as if it has been received (see 9.5.6).

10.2 Node based data structures

10.2.1 Unique ID

Each SSA node shall have a Unique ID that is globally unique. A UNIQUE ID field is an 8-byte field, comprised of two reserved bytes (containing zero) followed by a 6-byte IEEE Universal Address. The first 3 bytes of the IEEE Universal Address are an Organizationally Unique Identifier (OUI) that is assigned by IEEE³. The last 3 bytes are a product identifier that is administered by the OUI assignee.

As an example, the Organizationally Unique Identifier AC-DE-48 and the product identifier 000080h shall generate the Unique ID 0000.ACDE.4800.0080h.

The UNIQUE ID defined in Table 17 is used in the configuration process to detect cyclic paths and with the QUIESCE SMS to handle certain link errors. The UNIQUE ID is used to identify a particular node when the Web topology changes.

Table 17 - UNIQUE ID field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	reserved							
1	reserved							
2	IEEE ORGANIZATIONALLY UNIQUE IDENTIFIER							
3	IEEE ORGANIZATIONALLY UNIQUE IDENTIFIER							
4	IEEE ORGANIZATIONALLY UNIQUE IDENTIFIER							
5	PRODUCT IDENTIFIER							
6	PRODUCT IDENTIFIER							
7	PRODUCT IDENTIFIER							

10.2.2 Topology table

The Topology table shall be build and maintained by all Configurator nodes during the configuration process. The format of the Topology table is vendor specific, but at a minimum, the Topology table needs to have sufficient information to generate the Configuration table.

10.2.3 Configuration table

In the Configuration process (see 10.4) each Configurator node shall construct and maintain a Configuration table by walking the Web one node at a time. The table is essentially a list of all the nodes with pointers to represent the physical links.

The Master shall also issue a CONFIGURE PORT SMS to each port on every other node in the Web, including ports that are not currently Operational (see 12.2.3).

³To obtain an Organizationally Unique Identifier (OUI) contact: Registration Authority Committee, The Institute of Electrical and Electronic Engineers, Inc., 445 Hoes Lane, Piscataway NJ 08855-1331, USA, (908) 562-3812.

The format of the Configuration table is vendor specific. However, the Configuration table shall contain the information defined in Table 18 either directly or indirectly.

Table 18 - Configuration table entry

Field	Description
UNIQUE ID	The Unique ID of the remote node.
TOTAL PORTS	The total number of ports implemented by the remote node. This indicates the length of Pointer array below.
MASTER PRIORITY	A 3-bit field indicating the priority of the remote node for becoming the Web Master. Zero indicates that the node is not capable of functioning as a Master (that is, that the node is a Responder node). Non-zero values are the node s priority for becoming Web Master, with 1 the lowest priority and 7 the highest.
PATH ADDRESS	The address of the primary path to the node. Together with local port this provides the information needed to address the remote node rapidly.
LOCAL PORT	The number of the Configurator node port, for the primary path to the remote node.
RETURN PATH ID	The RETURN PATH ID of the primary path to the node.
RETURN PATH	The return path address from the node to the Configurator node. This is a required field in most SMSs that the Configurator node sends to the node. Saving it in the Configuration table optimizes performance.
SMS OUTSTANDING	A flag that is set when the Configurator node has an outstanding CONFIGURE PORT, MASTER ALERT, QUERY PORT, QUERY PROTOCOL, QUERY SWITCH, or QUIESCE SMS. The flag is reset when the SMS is no longer outstanding (an appropriate response is received, including an associated ASYNC ALERT SMS).

A Configuration table entry optionally contains other information provided by the upper-level protocol (e.g., the node type and serial number).

Availability and performance are improved if a Web is configured with more than one path between two particular nodes. In this case the choice of the primary path and any alternative paths is vendor specific.

If a node has a port with an OPERATIONAL FLAG cleared then the Configurator nodes do not attempt to walk that link during Configuration process (this shall produce an ASYNC ALERT SMS with an ALERT CODE value of ADDRESSED PORT NOT OPERATIONAL). If the port subsequently sets it s OPERATIONAL FLAG, each Configurator node receives a MASTER ALERT SMS with an ALERT CODE value of PORT NOW OPERATIONAL. Then each Configurator node shall walk the new link and adds nodes to its Configuration table until it encounters a single port node, another port with a cleared OPERATIONAL FLAG or it receives a Unique ID that is already known.

A Responder node shall not invoke the Configuration process or build a Configuration table.

10.2.4 Configurator table

In the configuration process each node shall construct a Configurator table from the information received in QUERY NODE SMSs. All nodes construct a Configurator table. The table associates a Return Path ID with Unique ID, Return Path and port. The Configurator table is used by a node when processing SMSs that use the RETURN PATH ID fields to associate the Return Path ID to a Return Path and port. The Configurator table is built when processing QUERY NODE SMSs, and modified when processing QUIESCE SMS or DELETE RETURN PATH ID SMSs.

The format of the Configurator table is vendor specific, however the Configurator table shall contain the information shown in Table 19 either directly or indirectly (i.e., the Return Path ID may be an index into the Configurator table). The Configurator table optionally contains information needed by the upper level protocols.

Table 19 - Configurator table entry

Field	Description
RETURN PATH ID ¹	The Return Path ID assigned to the referenced Unique ID, Port, and Return Path.
UNIQUE ID	The Unique ID of a Configurator node.
PORT	The node port to be used with the Return Path associated with the Return Path ID.
RETURN PATH	The Return Path to be used with the Port associated with the Return Path ID.
STATUS	The STATUS of the Configurator table entry as defined in Table 20.
NOTE ¹ - The Return Path ID field may be implied by indirect addressing of the table (i.e., The entry number is the same as the Return Path ID).	

Table 20 - Configurator table entry STATUS field

Value	Description
FREE	The Return Path ID is free for assignment and the associated UNIQUE ID, RETURN PATH and PORT elements are invalid.
HELD	The Return Path ID is held for the UNIQUE ID and the PORT and RETURN PATH fields are invalid.
VALID	The Return Path ID is assigned to the UNIQUE ID and uses the PORT and RETURN PATH. The Configurator table entry has changed since the last QUERY REGISTRATION SMS.
REPORTED	The Return Path ID is assigned to the UNIQUE ID and uses the PORT and RETURN PATH. The Configurator table entry has not changed since the last QUERY REGISTRATION SMS.

When a QUERY NODE SMS is received with the DR bit cleared, the node shall scan the Configurator table for an entry that matches all the following criteria:

- a) the STATUS field has a value of VALID or REPORTED;
- b) the UNIQUE ID field of the QUERY NODE SMS matches the UNIQUE ID element of the Configurator table entry;
- c) the RETURN PATH field of the QUERY NODE SMS matches the RETURN PATH element of the Configurator table entry;
- d) the port the QUERY NODE SMS was received on matches the PORT element of the Configurator table entry.

If a matching entry is found, then a QUERY NODE REPLY SMS is generated and the RETURN PATH ID field value from the Configurator table is copied into the RETURN PATH ID field of the SMS.

If no matching entry is found, then a search is made for an entry with a STATUS value of HELD and a matching UNIQUE ID field value. If no HELD entry match is found, then the first entry with a STATUS value of FREE is used. If a HELD or FREE entry is found, then the associated fields of the Configurator table entry are changed as follows:

- a) the PORT field is set to the port on which the QUERY NODE SMS was received;
- b) the STATUS field value is set to VALID
- c) the RETURN PATH field is set to the RETURN PATH field of the QUERY NODE SMS;
- d) the UNIQUE ID field is set to the UNIQUE ID field of the QUERY NODE SMS.

When a node executes a QUIESCE SMS it scans the Configurator table for all entries that have a UNIQUE ID element that matches the UNIQUE ID field of the QUIESCE SMS. If the RCTE bit is set in the QUIESCE SMS, then the STATUS field for all matching entries is set to FREE. If the RCTE bit is cleared in the QUIESCE SMS, then the STATUS field of the first matching entry is set to HELD, and the STATUS field of any other matching entries is set to FREE. Keeping one of the entries as HELD for that Unique ID guarantees that initiator may register at least one path. This quiesce action is performed when a previously active path from that Configurator node has been disconnected. The quiesced Configurator node is required to issue a QUERY NODE SMS to register over each path that it intends to continue using. This ensures that any unused RETURN PATHS are removed from the Configurator table.

The Configurator table is cleared (i.e., the STATUS fields of all entries are set to a value of FREE) when a node executes a Total Reset or Absolute Reset Control frame. When one of the links to the node becomes Operational again each Configurator node receives a MASTER ALERT SMS and performs the Configuration process which results in the Configurator table being rebuilt.

NOTE 8 - It is recommended that Responder nodes provide space for at least 64 entries in the Configurator table. If there is no space to make an entry when a node receives a QUERY NODE SMS then the node sets the ITF bit in the QUERY NODE REPLY SMS.

10.2.5 REFLECTION FLAG

The REFLECTION FLAG of a node, causes all ports on that node to either forward or reflect SAT and SAT characters as defined in 8.5.2.

The REFLECTION FLAG is set if any of the following conditions occur:

- a) the node is a single port node or switch node;
- b) one of the ports of a dual port node becomes not Operational;
- c) or a dual port node processes a CONFIGURE PORT SMS with the REFLECT bit set.

The REFLECTION FLAG is cleared when a dual port node with the OPERATIONAL FLAG set for both ports processes a CONFIGURE PORT SMS with the REFLECT bit cleared.

10.3 Beginning Communication process

A port shall enter the Disabled state and Privileged mode after the following conditions occur:

- a) POST;
- b) a Total Reset Control frame is processed;
- c) an Absolute Reset Control frame is processed;
- d) the Link ERP process exits with an ALERT CODE value;
- e) or a CONFIGURE PORT SMS is received with a non-zero NEGOTIATE SPEED field.

The port shall use the following procedure to begin communication in the order shown:

- a) If the line fault detector indicates a line fault, then wait until there is no line fault. A line fault indicates that the remote node is not powered-on or that the cable is disconnected.
- b) Invoke the Speed Negotiation process.
- c) Set the port s OPERATIONAL FLAG.
- d) Wait for at least 10 FLAG characters to be sent.
- e) Ensure that buffer space is available to receive at least 1 frame of any type. Then send an RR character pair and decrement the RR PENDING FLAG (see 9.4.3.).
- f) Wait to receive an RR character pair from the remote port. Then decrement the WAITING FOR RR FLAG.

Privileged frames may now be received and transmitted.

10.4 Speed Negotiation process

The Speed Negotiation process is exited whenever synchronization is achieved and a FLAG character is received (when the port enters the Ready state). The speed Negotiation process may broken down into two versions, single speed and dual speed.

10.4.1 Single Speed

The Speed Negotiation process for single speed devices sends at least 200 DIS characters to allow the remote port to synchronize prior to potentially forwarding a Control frame and then sends FLAG characters indefinitely until synchronization is achieved and a FLAG character is decoded. The single speed case of speed negotiation is as follows in order:

- a) Transmit at least 200 DIS characters;
- b) Put the port into the Enabled state;

- c) Wait until the port enters the Ready state.

10.4.2 Dual Speed

The Speed Negotiation process for dual speed devices attempts to synchronize at 400 MHz for 400 μ sec, then at 200 MHz for 200 μ sec, then at 400 MHz for 100 μ sec, then at 200 MHz for 300 μ sec then repeats the sequence indefinitely. At each speed, the process sends at least 200 DIS characters to allow the remote port to synchronize prior to possibly forwarding a Control frame. There is a 10% tolerance to the numbers to allow firmware implementation rather than forcing hardware implementation.

NOTE 9 - Implementers should minimize the tolerance, since a large tolerance increases the chances of synchronizing to a slower than optimal speed.

The following process is repeated indefinitely in the order shown until character synchronization is achieved:

- a) Set the transmitter and receiver to operate at 400 MHz and place the port in the Disabled state. Set a timer to time-out at 400 μ s \pm 10%. Transmit at least 200 DIS characters.
- b) If character synchronization is achieved the port enters the Enabled state changing the DIS characters to FLAG characters. A control frame may now be forwarded. Then if a FLAG character is decoded prior to the timer expiring, the port enters the Ready state and exits the Speed Negotiation process.
- c) Set the transmitter and receiver to operate at 200 MHz and place the port in the Disabled state. Set a timer to time-out at 200 μ s \pm 10%. Transmit at least 200 DIS characters.
- d) If character synchronization is achieved the port enters the Enabled state changing the DIS characters to FLAG characters. A control frame may now be forwarded. Then if a FLAG character is decoded prior to the timer expiring, the port enters the Ready state and exits the Speed Negotiation process.
- e) Set the transmitter and receiver to operate at 400 MHz and place the port in the Disabled state. Set a timer to time-out at 100 μ s \pm 10%. Transmit at least 200 DIS characters.
- f) If character synchronization is achieved the port enters the Enabled state changing the DIS characters to FLAG characters. A control frame may now be forwarded. Then if a FLAG character is decoded prior to the timer expiring, the port enters the Ready state and exits the Speed Negotiation process.
- g) Set the transmitter and receiver to operate at 200 MHz and place the port in the Disabled state. Set a timer to time-out at 300 μ s \pm 10%. Transmit at least 200 DIS characters.
- h) If character synchronization is achieved the port enters the Enabled state changing the DIS characters to FLAG characters. A control frame may now be forwarded. Then if a FLAG character is decoded prior to the timer expiring, the port enters the Ready state and exits the Speed Negotiation process.
- i) Return to item a.

10.5 Configuration

Each Configurator node shall invoke a Configuration process to determine the other nodes that are present and their path address(es).

10.5.1 Configuration Process

The Configuration Process shall involve the following steps in the order shown:

- a) Walk the Web (see 10.5.2) to build a Topology table (see 10.2.2).
- b) Examine the Topology table to determine the Master. A Master shall be defined as the Configurator node with the largest value reported in the MASTER PRIORITY field. If more than one node has the same value for the MASTER PRIORITY field, then the node with the largest Unique ID shall be the Master.
- c) Construct a Configuration table from the information in the Topology table.
- d) Register with the nodes that it expects to perform I/O process with. If a node supports multiple protocols a QUERY PROTOCOL SMS shall be issued to determine the protocols supported.
- e) If the node is to be the Master (see 13.1), then it shall internally places all of its operational ports in Normal mode.
- f) The Master shall send a CONFIGURE PORT SMS for each port to the nodes in the Topology table. Operational ports are placed in Normal mode, and non-operational ports are left unchanged.

- g) The Master shall send a MASTER ALERT SMS with an ALERT CODE value of BOTH PORTS OF LINK ARE IN NORMAL MODE for each link that has both ports in Normal mode. After configuring every port on a switch, the Master shall send a MASTER ALERT SMS with an ALERT CODE value of ALL OPERATIONAL PORTS OF ON THE SPECIFIED NODE ARE IN NORMAL MODE.
- h) When all operational ports of the Web have been placed in Normal mode, the Master shall send a MASTER ALERT SMS with an ALERT CODE value of ALL OPERATIONAL PORTS ON ALL NODES IN THE WEB ARE IN NORMAL MODE.

10.5.2 Web Walking Process

The Base node shall be set to the Configurator node performing the Web Walking process. All ports on the Configurator node shall be marked as unexplored.

The term Base node is defined as the node from which that the Configurator walks a string or loop. The Web Walking process is defined in the following steps in the order shown, and shall be performed until all ports are explored:

- a) The Configurator node selects a port on the base node and marks it as explored. The Web walking process is repeated for each unexplored operational port on the base node.
- b) The Configurator node shall send a QUERY NODE SMS with the DR bit set to the destination node incrementing the path address by one until a QUERY NODE REPLY is received reporting:
 - 1) only one port operational indicating the end of a string has been reached;
 - 2) a Unique ID that already exists in the Topology table indicating a loop has been established;
 - 3) or the TOTAL OTHER PORTS field value exceeds two indicating that a switch has been discovered.
- c) The Configurator node may send a QUERY PROTOCOL SMS and QUERY PORT SMSs as desired.
- d) If a switch is discovered the Configurator node shall send a QUERY SWITCH SMS to determine the number of operational ports on the switch. The base node shall be set to that switch node and the Web Walking process shall be recursively called. The port indicated in the QUERY NODE REPLY SMS shall be marked as explored.

10.6 Resets

This clause defines the various types of reset that are associated with SSA-TL2. ULPs that use SSA-TL2 may have reset functions of their own. Any ULP reset shall not affect any other ULP.

10.6.1 Link reset

The Link Reset Control frame is used during the Link ERP process to recover from a link error (see 11.1).

The Link Reset frame is confined to a single link. It shall never be forwarded by a router from one link to another. Therefore, the STATUS field contains the Link Status Byte that is used by the Link ERP process (see 11.1.2).

Every port shall be able to originate and invoke a Link Reset frame. A Link Reset frame shall only be originated by a port in the Check state, regardless of the WAITING FOR RR flag state.

Link Reset frames shall be acknowledged but are not subject to pacing or sequence numbers. The transmitter shall not inspect or change the WAITING FOR RR FLAG or the TRANSMIT SEQUENCE COUNTER. The receiver shall not inspect or change the RR PENDING FLAG or the RECEIVE SEQUENCE COUNTER

The receiver shall recognize a Link Reset whenever it is in the Ready or Check states, regardless of whether its inbound buffers are full. If the frame is valid the port shall return an ACK character pair and enter the Check state if it has not already done so (the ACK character pair is useful for fault isolation since it indicates that the link is Operational even if the node processor has hung). The contents of the frame buffers shall not be changed by the receipt of a Link Reset since they are needed by the Link ERP process. The port mode is not affected unless the Link ERP process exits with an ALERT CODE value.

10.6.2 Total Reset or Absolute Reset

A Total Reset frame resets the transport layer and notifies any registered ULPs of the reset. Hardware not directly under the control of the transport layer is not affected (i.e., the status of the spindle motor of a disk drive).

The destination node shall process a Total Reset Control frame as follows:

- a) perform optional resets and optional POST;
- b) clear the OPERATIONAL FLAG;
- c) all ports are set to the Disabled state and Privileged mode;
- d) the frame buffers shall be initialized;
- e) the TRANSMIT POINTER and RETRY POINTER shall be cleared;
- f) the Configuration table in a Configutor node shall be cleared;
- g) the Configutor table shall be cleared;
- h) the Port table shall be cleared;
- i) the Queued Asynchronous Alert table shall be cleared;
- j) the Asynchronous Alert Address table shall be cleared;
- k) any pending Asynchronous Alerts shall be cleared;
- l) the node shall initialize the following items for each port (see 12.2.3):
 - 1) A QUOTA = 1;
 - 2) B QUOTA = 4;
 - 3) REFLECT = 0b;
 - 4) EUDC = 0b;
 - 5) LINK ERP COUNTER = 0;
 - 6) and ALARM THRESHOLD = 10.
- m) the SSA-TL2 pool of TAG values is freed;
- n) the SSA-TL2 pool of channel values is freed;
- o) the Beginning Communication process shall be invoked within 25 ms of the receipt of a Total Reset frame.

Absolute Reset shall reset the transport layer and all ULPs, duplicating the effect of a power on reset.

When a Configutor node receives a Total Reset or Absolute Reset frames , it shall start a 10 second time-out and shall wait to receive a MASTER ALERT SMS with an ALERT CODE value of RECONFIGURATION REQUIRED. The Configutor node shall wait to configure the Web until after the time-out expires or the MASTER ALERT SMS is received. This is different from when a Configutor node shall configure the Web immediately following a power-on reset.

All Configutor nodes shall be capable of originating Total Reset, Reserved or Absolute Reset Control frames. The Total Reset or Absolute Reset frames shall not be subject to pacing, acknowledgments or sequence numbers. The transmitter shall sends the Total or Absolute Reset frame immediately after it finishes the current frame, if any. The receiver does not inspect or change the RR PENDING FLAG, ACK PENDING FLAG or RECEIVER SEQUENCE NUMBER.

A transmitter shall send a minimum of 200 FLAG characters prior to originating a Total Reset or Absolute Reset frame while in the Enabled State to allow the remote port time to synchronize if necessary. A transmitter shall be able to send the Total Reset or Absolute Reset frames when the port is in the Enabled or Ready state. The receiver shall recognize the Total Reset or Absolute Reset when the port has achieved character synchronization in any state, regardless of whether its inbound frame buffers are full.

The router in dual port or switch nodes shall be able to store and forward at least one Total Reset, Reserved or Absolute Reset Control frame. Since Control frames are not paced, dedicated storage is required if the frame buffers for Privileged, Reserved and Application frames are occupied. If another Total Reset, Reserved or Absolute Reset Control frame is received while the dedicated storage is in use then the node may discard it. If the output port is in the Disabled or Check state the Total Reset or Absolute Reset frame shall be held until it

enters the Enabled or Ready state. Consequently a Total Reset or Absolute Reset frame shall not be sent during the Link ERP process. The router shall be able to forward a Total Reset or Absolute Reset frame even when the output ports OPERATIONAL FLAG is cleared.

NOTE 10 - Any reset performed by a ULP shall not affect transport layer functions in such a way as to affect other ULPs running in the same node.

10.6.3 Power On Reset

Power on Reset is an internal function that resets the entire node at power-on, after a catastrophic internal error or when a node is disconnected from all other nodes.

When a node invokes a Power-on reset, it shall initialize the SSA-TL2 functions as follows:

- a) All ports shall enter Wrap mode;
- b) The node shall invoke an Absolute Reset.

10.6.4 Web Reset process

The Web Reset process may be used by the Master when communication over the Web is blocked and the cause cannot be identified.

The Master prepares to invoke a Web Reset process as follows in the order shown:

- a) The Master shall clear its Configuration table and shall make an entry for itself;
- b) The Master shall place each of its ports in the Disabled state and shall clear its OPERATIONAL FLAG. The Master shall execute the Beginning Communication process (see 10.3) for each of its ports.

The Master shall walk the entire Web resetting each node except itself. To avoid ambiguity, the following terminology is used:

- a) Current node is the node that is about to be reset;
- b) Previous port is the port nearer to the Master on the final link in the selected path from the Master to the Current node;
- c) Base node is the node from which the Master walks a string or loop.

For Webs that include switches, the Master walks one level of switching at a time. Initially, the Master is the Base node for resetting all nodes that are accessible without going through a switch node. Then, the Master uses each switch it discovers as the Base node for resetting all nodes that are accessible via a single port on each switch, and so on.

Initially, the Master marks all ports on each Base node as unselected. A Base node port becomes Selected when the Master uses it as the origin for walking a string or loop or when the Master encounters it as the last node of a string or loop.

The process to reset a loop or all of the strings connected to a particular Base node is as follows in the order shown:

- a) The Master selects an Unselected port on the Base node as the new Previous port and marks it as selected.
- b) The Master issues a CONFIGURE PORT SMS (or the internal equivalent) to Previous port. This ensures that the Master receives subsequent Asynchronous Alerts from the Previous port.
- c) The Master sends a QUERY NODE SMS to the Current node and starts a 125 ms timer.
- d) If the Master does not receive a QUERY NODE REPLY SMS before the timer expires, it continues with step e. Otherwise, if the Master receives a QUERY NODE REPLY SMS, the Master shall compare the UNIQUE ID field in the QUERY NODE REPLY SMS with each entry in its Configuration table. If the Unique ID is in the Configuration table then the Master shall mark the port indicated by the Current port field in QUERY NODE REPLY SMS as Selected and shall skip to step i; the Master has returned on a loop or a cyclic path to itself or a switch that it has already reset. If the Unique ID is not in the Configuration table, the Master shall continue with step e; since the Current node has not yet been reset.

- e) The Master sends a Total Reset to the Current node, starts a 125 ms timer, and waits to receive an ASYNC ALERT SMS with an ALERT CODE value of PORT NOW OPERATIONAL from Previous port. Except for issuing an ASYNC REPLY SMS, the Master shall ignore any other ASYNC ALERT SMSs (for example the Total Reset may cause the Previous port to initially invoke the Asynchronous Alert process with an ALERT CODE value of REMOTE PORT DISABLED).
- f) If the Master does not receive an ASYNC ALERT SMS with an ALERT CODE value of PORT NOW OPERATIONAL from Previous port before the 125 ms timer expires then the Master shall send an Absolute Reset to the Current node, assume Current node is not operational and skip to step i.
- g) If the Master does not receive a QUERY NODE REPLY SMS before the timer expires, then the Master shall skip to step i; the Current node shall be assumed to be not operational. Otherwise the Master shall make an entry in its Configuration table for the Current node.
- h) If the Current node is a single port node, the Master shall continue with step i. If the Current node is a dual port node, the Master updates the Previous port to refer to the port on Current node that is further from the Master, identifies the next node from this port as the new Current node and returns to step b. If the Current node is a switch, the Master shall mark the port indicated by the PORT field in QUERY NODE REPLY SMS as Selected and shall continue with step i.
- i) If there is an Unselected port remaining on the Base node, the Master returns to step a. Otherwise the walk from the current Base node is complete and the Master selects a new Base node if one exists.

When the Web Reset process is complete the Master shall clear its Configuration table and shall invoke the Configuration process described in clause 10.4. The Master shall also send a MASTER ALERT SMS with an ALERT CODE value of RECONFIGURATION REQUIRED to each other Configurator node. This shall cause each other Configurator node to clear its Configuration table and invoke the Configuration process.

11 Error handling

This clause defines the strategy for recovering errors in a complex SSA Web with multiple Configurator nodes. The strategy ensures data integrity and it minimizes the impact to other operations. For example, an error does not result in bad data being written to a device. An error affects only the commands and nodes currently using the failing link or node.

11.1 Link ERP

SSA-TL2 defines a Link Error Recovery Procedure (Link ERP) to recover link errors at the frame level. This has the following benefits:

- a) The upper-level protocol is simplified since frame error recovery is transparent if it is successful;
- b) There is normally no need to terminate any commands when a transient error occurs;
- c) There is no uncertainty about the state of the remote node.

If the Link ERP process determines that a transmission error occurred then it attempts to recover from the error itself. If recovery is successful the Link ERP process terminates and the upper-level protocol continues unaware of the error.

The port that detects the Link Error invokes the Link ERP process. During the Link ERP process, a Link Reset frame shall be sent to the remote node, that shall trigger the remote node to invoke the Link ERP process (if it did not also detect an error and invoke its own Link ERP process). The Link ERP process is coordinated between the two ports so that they remain synchronized.

The Link ERP process cannot recover from some errors transparently (e.g., hardware errors or permanent line faults). The Link ERP process has been designed so that both nodes always recognize an unrecoverable error and remain synchronized. In these cases the Link ERP exits and attempts to alert the Master via the Asynchronous Alert process (see 11.3).

The basic principles of the Link ERP process are as follows:

- a) Only the failing link invokes the Link ERP process. Other links in the Web are not involved.
- b) The Link ERP process does not recover Control frames.

- c) In normal operation the transmitter does not discard a Privileged, Application, or Reserved frame until it has received an ACK character pair. This indicates that the frame has been received correctly by the remote port. Therefore when an error occurs the affected frame(s) are still available for re-transmission without reference to the upper-level protocol.
- d) When an error is detected both ports enter the Check state, invoke the Link ERP process and exchange status by means of Link Resets.
- e) Recovery is invoked separately for each line. Each port is responsible for recovering frames that were lost on its outbound line. If the transmitter starts sending another frame before it receives an ACK character pair and a recoverable error occurs, 1 or 2 frames are retransmitted.
- f) Before restarting communication, the Link ERP process forces the port into the Disabled state. This synchronizes the Link ERP process in both nodes and allows an orderly restart with the same mechanisms that are used at power-on.
- g) The link protocol and ERP are designed to eliminate the loss or duplication of frames when an error occurs.

11.1.1 Link Errors

Except where explicitly stated the following errors are only detected when a port is in the Ready state before the error. Errors are ignored if a port is not in the Ready state.

When an error is detected, the port enters the Check state and interrupts the node processor. No further Privileged, Reserved or Application frames are accepted or acknowledged until the node returns to the Ready state.

11.1.1.1 Hardware error

The Hardware error indicates when a port detects an internal hardware error (e.g., a parity check).

11.1.1.2 Line Fault

The Line Fault error indicates when the output of the line fault detector is active and the port is not in the Disabled state or Wrap mode (e.g., the cable is open, short circuited or the remote node is powered off).

11.1.1.3 ACK Time-out

The ACK Time-out error indicates that the source port has not received an ACK character pair within the specified time after transmitting the trailing FLAG character of a frame other than a Total Reset, Reserved or Absolute Reset Control frame.

11.1.1.4 LOSS OF SYNCHRONIZATION ERROR

The LOSS OF SYNCHRONIZATION ERROR indicates that the clock recovery circuits in the receiver detected a synchronization error in the Ready state.

11.1.1.5 CODE VIOLATION ERROR

The CODE VIOLATION ERROR indicates that the receiver has not detected a LOSS OF SYNCHRONIZATION ERROR, and has either decoded a character that is not a data or Special character or has decoded a character that causes a disparity error.

11.1.1.6 PROTOCOL ERROR

The PROTOCOL ERROR indicates the receiver has neither detected a LOSS OF SYNCHRONIZATION ERROR nor has detected a CODE VIOLATION ERROR, and a port receives an incorrect sequence of valid characters as follows:

- a) a frame with less than 6 data characters between 2 FLAG characters (e.g., noise corrupting or manufacturing a FLAG character);
NOTE 11 - It is recommended that a short, and therefore invalid, Privileged, Reserved or Application frame with less than 7 data characters between 2 FLAG characters also be identified as a PROTOCOL ERROR.
- b) a Privileged, Reserved or Application frame and no buffer is available, (i.e., when the RR PENDING FLAG is non-zero);
- c) an unexpected RR character pair (i.e., when the WAITING FOR RR FLAG is zero);
- d) an isolated RR character, or 3 consecutive RR characters. One half of a link hangs if an RR character pair is lost without any errors being detected, (e.g., if both RR characters are changed to FLAG characters while the link is idle). This is extremely unlikely but the Master Alive process shall detect and correct this case;
- e) an unexpected ACK character pair (i.e., when the WAITING FOR ACK FLAG is zero);
- f) an isolated ACK character, or 3 consecutive ACK characters. If an ACK character pair is corrupted then the transmitter also detects an ACK Time-out error;
- g) a NUL character with no intervening data character since the last FLAG character:
- h) an ABORT character with no intervening data character since the last FLAG character:
- i) an ABORT character that is not immediately followed by a FLAG character;
- j) an isolated DIS character.

11.1.1.7 CRC ERROR

The CRC ERROR indicates that a frame was received that was not aborted and none of the receiver errors with a lower value than CRC ERROR has occurred (see Table 22).

11.1.1.8 SEQUENCE ERROR

The SEQUENCE ERROR indicates that all of the following conditions are true:

- a) the FSN field of a frame does not equal the port s RSC;
- b) the frame has not been aborted;
- c) the frame is not a Control frame;
- d) none of the receiver errors has occurred with a lower value than SEQUENCE ERROR (see Table 22).

11.1.1.9 FRAME REJECT ERROR

The FRAME REJECT ERROR indicates that a frame is received correctly with none of the receiver errors above, the frame has not been aborted, and the frame is unacceptable for one of the following reasons:

- a) The frame contains more than 139 data characters. The receiver shall continue to accumulate the CRC until the trailing FLAG, since a CRC ERROR Link ERP exit takes precedence over a FRAME REJECT Link ERP exit.
- b) The destination port determines that the CONTROL field contains a Reserved FRAME TYPE or Reserved RESET TYPE field (see 8.1).
- c) The first byte of the Path component is invalid (see 9.2.1, 9.2.2 and 9.2.3).
- d) The destination port determines that the Channel address is invalid or the destination port cannot accept a frame on the specified Channel (see 8.5.4).
- e) The length of the DATA field in a Control frame is not zero.
- f) The destination port determines that the length of the DATA field in an SMS frame exceeds 32 bytes.
- g) The destination port determines that the length of the DATA field in a frame exceeds 128 bytes.
- h) The length of the DATA field in a Data frame is not an integer multiple of 8 bytes, the frame does not contain the last byte of the transfer and the destination port cannot accept it.

11.1.2 Link Status Byte

During error recovery the Link ERP process in each node builds a Link Status Byte as defined in Table 21 and sends it to the other node in the STATUS field of a Link Reset frame.

Table 21 - Link Status Byte

Bit 7	6	5	4	3	2	1	Bit 0
HW	LF	ACK	RECEIVER ERRORS			RSC	

The HW bit when set indicates that the port detected an internal hardware error.

The LF bit when set indicates that the port detected a line fault. It is provided for diagnostic information only and it is not referenced by the Link ERP process in the remote node.

The ACK bit when set indicates that the transmitter timed-out while waiting for an ACK character pair. It is provided for diagnostic information only and it is not referenced by the Link ERP process in the remote node.

The RECEIVER ERRORS field is defined in Table 22.

Table 22 - Receiver errors

Value	Description
000b	NO ERROR
001b	LOSS OF SYNCHRONIZATION ERROR
010b	CODE VIOLATION ERROR
011b	PROTOCOL ERROR
100b	CRC ERROR
101b	SEQUENCE ERROR
110b	FRAME REJECT ERROR
111b	reserved

When two or more errors occur simultaneously the lowest RECEIVER ERROR field value is reported (i.e., SEQUENCE ERROR is lower priority than CRC ERROR).

The RSC field specifies the RECEIVE SEQUENCE COUNTER for the next Privileged, Reserved or Application frame that was acknowledged by the port. It is needed by the Link ERP process in the remote node.

11.1.3 Link ERP process definition

To facilitate cross-referencing, the Link ERP exits are identified by an ALERT CODE value that corresponds to the ALERT CODE field used by the Asynchronous Alert process.

The Link ERP process shall not be invoked from within itself (i.e., it is non-recursive). The implementation shall protect against the Link ERP process looping if there is a permanent error.

NOTE 12 - The following is an example of one method that is used to protect against the Link ERP process from looping. Each invocation of the Link ERP process increments a RETRY COUNTER that is reset to zero periodically by a timer. If the number of retries in one period of the timer exceeds a maximum value then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, it places the port in the Disabled state, and it exits with an ALERT CODE value of RETRY LIMIT EXCEEDED. This example also protects against excessive use of the ERP in the event of severe external noise.

The port that detects the error enters the Check state and invokes its Link ERP process. Any Link ERP exit shall terminate the Link ERP process and invoke the Asynchronous Alert process as defined in 11.3. The Link ERP process is as follows in order:

- a) The Link Status Byte is built. If the port detects a line fault that persists for more than 1 ms the port enters the Disabled state, and Link ERP process exits with an ALERT CODE value of PERMANENT LINE FAULT.

- b) If the port is not receiving characters, the port enters the Disabled state, and the Link ERP process exits with an ALERT CODE value of NO CHARACTERS RECEIVED.
- c) If the port is detecting DIS characters the port enters the Disabled state and the Link ERP process exits with an ALERT CODE value of REMOTE PORT DISABLED.
- d) Either a Link Reset frame (normal mode) or an Extended Link Reset frame (Extended Distance Option) is transmitted. If no acknowledgment is received the Link Reset or Extended Link Reset frame is retransmitted once. If no acknowledgment is received to the retransmitted Link Reset or Extended Link Reset frame, then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, sets the port to the Disabled state, and exits with an ALERT CODE value of LINK RESET FAILED.
- e) If the LINK RESET RECEIVED FLAG is cleared and a Link Reset frame has not been received within 5 ms, the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, places the port in the Disabled state, and exits with an ALERT CODE value of LINK RESET FAILED.
- f) If a hardware error is detected, the Link ERP process sets the port to the Disabled state and exits with an ALERT CODE value of HARDWARE ERROR.
- g) If FRAME REJECT ERROR is indicated by the port's Link Status Byte, the Link ERP process sets the port to the Disabled state, and exits with an ALERT CODE value of FRAME REJECT.
- h) The number of outbound frames with an acknowledgment outstanding is calculated. When using strict cyclic buffers this is obtained as follows (see 10.1.2):

$$Q = (\text{TRANSMIT POINTER} - \text{RETRY POINTER}) \text{ modulo } N$$

where N is the number of buffers in use by the transmitter. Q is integral number of frames up to the number of buffers.

The number of outbound frames that the local port is expecting an acknowledgment for but have not been received by the remote port is also calculated, where local refers to the ports TRANSMIT SEQUENCE COUNTER and remote refers to the RSC field received in the Link Status Byte.

$$P = (\text{local TRANSMIT SEQUENCE COUNTER} - \text{remote RSC}) \text{ modulo } M$$

where M is either 4 (non-Extended Distance Option) or 256 (Extended Distance Option) and where P is less than or equal to Q .

If either $Q > \text{number of buffers}$ or $P > Q$ then the Link ERP process waits 25 ms to ensure that the remote port recognizes an unrecoverable error, sets the port to the Disabled state and it exits with an ALERT CODE value of INVALID RETRY STATUS.

- i) The Link ERP process arranges to resend the number of frames indicated by subtracting P from its TRANSMIT POINTER, modulo N . Those outbound buffers that do not need to be retransmitted shall be discarded using the following algorithm.

```

Do while RETRY POINTER  $\neq$  TRANSMIT POINTER;
    Discard buffer at RETRY POINTER;
    Increment RETRY POINTER modulo  $N$ ;
End;

```

The frame buffers and TRANSMIT POINTER are set to retransmit the proper number of frames after the Link ERP re-establishes communication following the port being set to the Disabled state.

- j) If the port has received a frame containing any of the errors listed in Table 18, then the appropriate inbound buffer contents shall be discarded. The ERP places the port in the Disabled state. The ERP waits until the receiver detects DIS characters. This is required to synchronize the local and remote ports Link ERP process's and prevent the transmitter sending RR while the remote port is still in the Check state.

If the receiver does not detect DIS characters within 5 ms after the local port is disabled then the ERP exits with an ALERT CODE value of TIME-OUT WAITING FOR DISABLED STATE.

- k) The port is set to the Enabled state.

If the receiver detects a FLAG character within 5 ms of placing the port in the Enabled state, the port is set to the Ready state, otherwise the Link ERP process exits with a meaning code of TIME-OUT WAITING FOR READY STATE.

- l) The ERP terminates successfully.

11.1.3.1 Timing restrictions

To avoid false time-outs the implementation shall ensure that the following paths in the Link ERP process do not exceed 5 ms each:

- a) from receiving a Link Reset frame to sending a Link Reset frame;
- b) from the later of sending or receiving a Link Reset frame to placing the local port in the Disabled state;
- c) from the later of detecting DIS characters or placing the local port in the Disabled state to placing the local port in the Enabled state.

11.1.3.2 ERP exits

If the Link ERP process exits unsuccessfully then each node shall take the following actions in order:

- a) Clear the port's OPERATIONAL FLAG.
- b) If the port is capable of the Extended Distance Option, then it disables its Extended Distance Option and sets the ACK time-out to between 100 and 200 μ sec.
- c) Set the port to Privileged mode. The transmitter discards Application frames. This avoids frames backing up and blocking other traffic. It also prevents any further data from being transferred through the affected port.
- d) Invoke the Asynchronous Alert process (see 11.3) with the ALERT CODE specified in the Link ERP exit.
- e) Restart the affected port, as described in 10.3.

11.1.3.3 Example

The operation of the Link ERP process is illustrated in Table 23. Port 1 sends 2 frames back-to-back. Port 2 receives the first frame correctly but the ACK character pair is corrupted by noise. Port 1 then detects a code violation and aborts the second frame. After the Link ERP process completes Port 1 retransmits the second frame. To illustrate the operation of the TRANSMIT POINTER (TP) and the RETRY POINTER (RP) it is assumed that Port 1 has at least 3 transmit buffers.

Table 23 - Recovery from corrupted ACK character pair

Port 1	Line →	← Line	Port 2
TP=0, RP=0, TSC=0, RSC=0			TP=0, RP=0, TSC=0, RSC=0
Tx frame 1	Frame 1		
	"	RR character pair	Tx RR character pair
	"		
	"		
TSC=1, TP=1			Good CRC
Tx frame 2	Frame 2		RSC=1
	"	** Error **	Tx ACK character pair
Code violation	"	RR character pair	Tx RR character pair
Enter Check state	"		
Invoke Link ERP process	ABORT		RSC=1 (still)
Wait for Tx complete			
Assemble LSB			
Tx Link Reset	Link reset		
Wait to Rx Link Reset		ACK character pair	Tx ACK character pair
"			Enter Check state
"			Invoke Link ERP process
"			Wait for Tx complete
"			Assemble LSB
"		Link reset	Tx Link Reset
Tx ACK character pair	ACK character pair		Q=0, P=0
Q=1, P=0			
Discard 1 frame			
RP =1			
Enter Disabled state	DIS	DIS	Enter Disabled state
TSC=0, RSC=0	"	"	TSC=0, RSC=0
Wait to detect DIS	"	"	Wait to detect DIS
Enter Enabled state	FLAG	FLAG	Enter Enabled state
Wait to Rx FLAG			Wait to Rx FLAG
Enter Ready state			Enter Ready state
Tx frame 2	Frame 2		
	"	RR character pair	Tx RR character pair
	"		
	"		Good CRC
TSC=1, TP=2			RSC=1
		ACK character pair	Tx ACK character pair
RP=2			

11.2 Invalid SMS field handling

This clause discusses invalid field handling of data fields within frames. It is assumed that the frame has been successfully received by the hardware. This discussion centers on the processing of the DATA field (i.e., the contents of the SMS).

The order of processing is critical and matches the order of the following sub clauses.

11.2.1 Unknown SMS code

An SMS received with an unknown SMS CODE cannot be parsed since it cannot be determined if the SMS has either a RETURN PATH or RETURN PATH ID field, and therefore shall invoke an Asynchronous Alert process with an ALERT CODE value of UNKNOWN SMS.

11.2.2 Unsupported SMS code

An SMS received with a known SMS CODE that is either not supported by the node, or has the incorrect FRAME TYPE shall invoke an Asynchronous Alert process with an ALERT CODE value of UNKNOWN SMS.

11.2.3 SMS too small

If the SMS length is shorter than specified under clause 11, an Asynchronous Alert process shall be invoked with an ALERT CODE value of SMS TOO SHORT.

It is not an error to receive an SMS longer than the specified length.

11.2.4 Invalid RETURN PATH ID

An Invalid RETURN PATH ID is a RETURN PATH ID not contained in the Configurator table. No RETURN PATH and port is associated with it and therefore, no response is sent to the specified Configurator node.

If the Responder node receives an SMS with an invalid RETURN PATH ID, an Asynchronous Alert process is invoked with an ALERT CODE value of INVALID RETURN PATH OR RETURN PATH ID. When the Master generates the MASTER ALERT SMS, this causes all Configurator nodes to issue a QUERY NODE SMS to the associated node to rebuild its Configurator table.

11.2.5 Invalid Tag

An SMS sent to a Configurator node with an invalid TAG, or a valid TAG with an improper response causes the Configurator node to ignore the SMS; and optionally log the error. In the case of data transfer SMSs with an invalid TAG, the protocol layer issues the appropriate upper level protocol response to abort the I/O process if possible (i.e., ABORT TAG SMS for SSA-S2P or SSA-S3P).

11.2.6 Invalid context

When an SMS is received out of context, it shall be discarded and the event shall be logged by a Configurator node or ignored by a Responder node. Examples of out of context SMSs include the following:

- a) DATA REPLY SMS when no associated DATA READY SMS was issued;
- b) DATA REQUEST SMS was received for an I/O process that did not need to write data;
- c) DATA READY SMS was received for an I/O process that did not need to read data;
- d) ASYNC REPLY SMS with a valid TAG value without an outstanding ASYNC ALERT SMS;
- e) QUERY PROTOCOL REPLY SMS with a valid TAG value not associated with an outstanding QUERY PROTOCOL SMS;
- f) QUERY PORT REPLY SMS with a valid TAG value not associated with an outstanding QUERY PORT SMS.

11.2.7 Non-zero reserved field or reserved code value

An otherwise valid SMS with a non-zero reserved field or a reserved code value is parsed and the RETURN PATH ID or RETURN PATH field is used to send the appropriate response (excluding Unknown or Invalid SMS codes). The response for each transport layer SMS is defined in Table 24.

Table 24 - Response to non-zero reserved fields or reserved code values

SMS Name	Resulting SMS	Notes
QUERY NODE	Normal QUERY NODE REPLY	Ignore reserved field/values
QUERY NODE REPLY	None	Ignore reserved field/values
CONFIGURE PORT	RESPONSE SMS with RETURN CODE of INVALID FIELD	
RESPONSE	Discard	System dependent retry
ASYNC ALERT	ASYNC REPLY SMS with RETURN CODE of INVALID FIELD	Generate MASTER ALERT
MASTER ALERT	RESPONSE SMS with RETURN CODE of INVALID FIELD	
QUIESCE	RESPONSE SMS with RETURN CODE of INVALID FIELD	
ASYNC REPLY	None	No resend needed, clear the ASYNC ALERT data.
QUERY PROTOCOL	RESPONSE SMS with RETURN CODE of INVALID FIELD	
QUERY PROTOCOL REPLY	Discard	System dependent retry
QUERY PORT	RESPONSE SMS with RETURN CODE of INVALID FIELD	
QUERY PORT REPLY	Discard	System dependent retry
DATA READY	Discard	Cancel I/O Process
DATA REPLY	RESPONSE SMS with RETURN CODE of INVALID FIELD	
DATA REQUEST	Discard	Cancel I/O Process

11.3 Asynchronous Alert process

The Asynchronous Alert process is given an ALERT CODE value, a port, and a frame (if applicable). All tests are based on the various tables associated with the Port parameter. The node then invokes one of the following options in order.

- a) If a QUERY NODE REPLY SMS has not been sent since the last power on, Total Reset or Absolute Reset then discard the data and exit the Asynchronous Alert process;
- b) Otherwise if the AA VALID flag is cleared, or the Queued Asynchronous Alert table is not empty, then queue the parameters in the Queued Asynchronous Alert table and exit the Asynchronous Alert process;
- c) Otherwise, queue the Asynchronous Alert data and invoke the De-queue Asynchronous Alert process with the Port parameter (see 11.4).

The receipt of the ASYNC ALERT SMS by the Master invokes the ASYNC ALERT SMS Handling process

11.4 De-queue Asynchronous Alert process

The De-queue Asynchronous Alert process is invoked by the Asynchronous Alert process for the first Asynchronous Alert following a valid CONFIGURE PORT SMS (AA VALID flag set in the Port table) or by the receipt of an ASYNC REPLY SMS that terminates an outstanding ASYNC ALERT SMS.

The De-queue Asynchronous Alert process is given a Port parameter. All tests are based on the various tables associated with the Port parameter. The De-queue Asynchronous Alert process is as follows in order:

- a) If the Queued Asynchronous Alert table is empty then exit the De-queue Asynchronous Alert process
- b) Build an ASYNC ALERT SMS using the port (PORT field), ALERT CODE value, and frame (CONTROL, CHANNEL, and FRAME DATA fields). The TAG field value and path for the ASYNC ALERT SMS are obtained from the AA TAG and AA RETURN path fields of the Port table respectively.
- c) Send the ASYNC ALERT SMS. If no ASYNC REPLY SMS with a matching TAG and ALERT CODE value is received within 500 ms, then resend the ASYNC ALERT SMS. Repeat this step indefinitely until either a matching ASYNC REPLY SMS is received, or the port is reset by a Power On, Total, or Absolute Reset.

11.5 ASYNC ALERT SMS Handling process

Asynchronous Alerts are used to report events that are outside the scope of the Link ERP process:

- a) a new link is connected to the Web;
- b) a router receives a frame that is addressed to a port that is not Operational;
- c) a node receives an invalid or unknown SMS;

- d) a node detects a protocol error;
- e) a node detects a critical environmental error;
- f) a node detects a non-critical warning error;
- g) a Link ERP process exit;
- h) a node detects a recovered Link ERP process threshold is exceeded;
- i) a change in the list of ULPs supported has occurred since the last QUERY NODE SMS or QUERY PROTOCOL SMS.

Upon receipt of an ASYNC ALERT SMS, the Master optionally logs the error. Any other action required to handle the different ALERT TYPE field values is defined in the following sub-clauses.

11.5.1 General ASYNC ALERT SMS Handling process

The receipt of an ASYNC ALERT SMS with the following TYPE field values shall invoke the general ASYNC ALERT SMS Handling process:

- a) ADDRESS ERROR is used when a frame is addressed to a port with a cleared OPERATIONAL flag. The port detecting the error remains in Normal mode.
- b) UNKNOWN SMS is used when the SMS is not recognized by the node.
- c) INVALID SMS is used when the SMS is deemed invalid.
- d) PROTOCOL ERROR is used when an SMS is received out of order.
- e) ENVIRONMENTAL ERROR is used when a critical system failure is detected. The handling of the error is system dependent, but often results in the suspension of all activity in the system.
- f) ERROR THRESHOLDS is used when a non-critical system degradation is detected by exceeding a predetermined number of recovered link errors. The handling of this error is system dependent, but often results in system activity continuing (with a possible rerouting of future activity) with the possible generation of a repair action.

The general ASYNC ALERT SMS Handling process is defined as follows in order:

- a) Build an ASYNC REPLY SMS using the TAG and ALERT CODE values from the ASYNC ALERT SMS.
- b) Send the ASYNC REPLY SMS using the ADDRESS field in the ASYNC ALERT Address table entry associated with the TAG field value of the ASYNC ALERT SMS.
- c) Build a MASTER ALERT SMS using the port (PORT field), ALERT CODE value, and frame (CONTROL, CHANNEL, and FRAME DATA fields).
- d) Send the MASTER ALERT SMS to all Configurator nodes using the PATH ADDRESS field of the Configuration table entries.
- e) Log the error.
- f) Wait for a RESPONSE SMS from each of the Configurator nodes for the MASTER ALERT SMS. If any Configurator node fails to return a RESPONSE SMS associated with the MASTER ALERT SMS within 1 second, the Web Reset process shall be invoked (see 10.6.4).

The receipt of the MASTER ALERT SMS invokes the Master Alert Handling process by the Configurator node.

11.5.2 PORT NOW OPERATIONAL Type specific handling process

When a new link is connected to the Web a port on one of the existing nodes sets its OPERATIONAL FLAG and sends an ASYNC ALERT SMS to the Master with an ALERT CODE of PORT NOW OPERATIONAL. The PORT NOW OPERATIONAL Type specific handling process is as follows in order:

- a) Perform the general ASYNC ALERT SMS Handling process (see 11.5.1).
- b) The Master shall walk the new link by issuing QUERY NODE SMSs to the new node(s). Each node shall return a QUERY NODE REPLY SMS.
- c) The Master shall issue a CONFIGURE PORT SMS specifying SET NORMAL MODE to those ports that are Operational on the new nodes that it has walked. Each port shall return a RESPONSE SMS.
- d) The Master shall issue a Master alert to each other Configurator node with an ALERT CODE value of BOTH PORTS OF LINK ARE IN NORMAL MODE for each new link whose ports are now in Normal mode. If

the Web is now a loop, the Master optionally issues a single MASTER ALERT SMS with an ALERT CODE value of ALL OPERATIONAL PORTS ON ALL NODES IN THE WEB ARE IN NORMAL MODE in place of the individual Master alerts with an ALERT CODE value of BOTH PORTS OF LINK ARE IN NORMAL MODE.

- e) After the other Configurator nodes have returned a RESPONSE SMS for each Master alert, no additional ASYNC ALERT SMSs with a error code of PORT NOW OPERATIONAL are processed until this procedure is completed.

11.5.3 WARNING Type specific handling process

The actions taken when a WARNING TYPE FIELD value are received are separated into three groups as follows:

- a) The WARNING Type MASTER PRIORITY CHANGED ALERT CODE Asynchronous Alert is generated when the MASTER PRIORITY level of a Configurator node changes. The actions taken are defined in 13.1.
- b) The WARNING Type SUPPORTED ULP LIST CHANGED ALERT CODE Asynchronous Alert is generated when a node changes the ULP that it supports. Any Configurator receiving this MASTER ALERT SMS may respond to this warning by initiating a QUERY NODE SMS and a QUERY PROTOCOL SMS, if needed, to the node causing the alert. The supported protocol list received is used to update the table data held in regards to that node. If a Configurator for any reason does not need the changed information, then the warning may be discarded with no action.
- c) The WARNING Type excluding MASTER PRIORITY CHANGED and SUPPORTED ULP LIST CHANGED ALERT CODE Asynchronous Alert types are generated when a non-critical system degradation is detected. The handling of this error is system dependent, but often results in system activity continuing with the generation of a repair action. The Master alerts all Configurator nodes who takes the appropriate system dependent action.

11.5.4 LINK ERP EXITS Type specific handling process

An ERP exit indicates that an unrecoverable link error has occurred. Initially the Master is informed by one or two ASYNC ALERT SMSs from the nodes that are connected by the failing link. The Master returns an ASYNC REPLY SMS and recovery proceeds as follows in order:

- a) Perform the general ASYNC ALERT SMS Handling process (see 11.5.1).
- b) If the ASYNC ALERT SMS has an ALERT CODE value of either LINK RESET FAILED or TIME-OUT WAITING FOR READY STATE following an ERP exit then the Master sends a Total Reset frame to the failing node.
- c) The Master invalidates the entries in the Configuration table for the affected link(s) and node(s). If the failing node recovers after the Total Reset then the remote node(s) invoke the Asynchronous Alert process with an ALERT CODE value of PORT NOW OPERATIONAL. This causes the affected links and nodes to be reconfigured.
- d) If the ASYNC ALERT SMS indicates any Link ERP process exit, the Master shall invalidate the entries in the Configuration table for the affected link(s).
- e) If the Master determines that a previous Configurator node has been removed from the Web by the error then it sends a QUIESCE SMS with the RCTE bit set to all remaining nodes on behalf of the missing Configurator node.
- f) If the link is still configured the Master sends a CONFIGURE PORT SMS with a MODE field value of SET NORMAL MODE to each port. The Master waits for the RESPONSE SMSs for CONFIGURE PORT SMS. The Master sends a MASTER ALERT SMS to each other Configurator node with an ALERT CODE value of BOTH PORTS OF A LINK ARE IN NORMAL MODE. Each Configurator node returns a RESPONSE SMS.

11.5.5 MASTER GENERATED ALERT Type specific handling

The MASTER GENERATED ASYNC ALERT TYPE field value shall only be generated by the Master, and therefore the Master never receives an ASYNC ALERT SMS with a MASTER GENERATED ALERT TYPE field. If the Master receives an ASYNC ALERT SMS with a MASTER GENERATED ALERT TYPE field, then the Master performs the following actions in order:

- a) Build an ASYNC REPLY SMS using the TAG and ALERT CODE values from the ASYNC ALERT SMS.
- b) Send the ASYNC REPLY SMS using the ADDRESS field in the ASYNC ALERT Address table entry associated with the TAG field value of the ASYNC ALERT SMS and the RETURN CODE field value of INVALID FIELD.
- c) Logs the error.

A MASTER ALERT SMS is not sent to the other Configurator nodes.

11.6 MASTER ALERT SMS Handling process

11.6.1 General MASTER ALERT SMS Handling process

The receipt of an MASTER ALERT SMS with the following TYPE field values shall invoke general MASTER ALERT SMS Handling process:

- a) UNKNOWN SMS causes vendor specific error handling;
- b) INVALID SMS causes vendor specific error handling;
- c) PROTOCOL ERROR causes vendor specific error handling;
- d) ENVIRONMENTAL ERROR causes vendor specific error handling;
- e) WARNING causes vendor specific error handling;
- f) ERROR THRESHOLDS causes vendor specific error handling.

The receipt of an MASTER ALERT SMS by the Configurator node invokes the MASTER ALERT SMS Handling process defined as follows in order:

- a) build a RESPONSE SMS using the TAG value from the MASTER ALERT SMS;
- b) send the RESPONSE SMS using the RETURN PATH value in the MASTER ALERT SMS as a frame ADDRESS field value;
- c) inform the upper layer protocol if appropriate;
- d) log the error.

11.6.2 PORT NOW OPERATIONAL type handling process

The PORT NOW OPERATIONAL Type specific handling process is as follows in order:

- a) Invoke the Configuration process starting at the node and port that is now operational;
- b) Perform the General MASTER ALERT SMS Handling process (see 11.6.1).

11.6.3 ADDRESS ERROR Type Handling process

The ADDRESS ERROR Type specific handling process is as follows in order:

- a) Terminates any affected I/O processes and the associated out-bound data transfers;
- b) Each Configurator node invalidates the entries in the Configuration table for the affected link(s) and node(s);
- c) Perform the General MASTER ALERT SMS Handling process (see 11.6.1);
- d) If there is an alternative path, reissue the affected commands.

11.6.4 LINK ERP EXITS Type Handling process

The LINK ERP EXITS Type specific handling process is as follows in order:

- a) Invalidate the entries in the Configuration table for the affected link(s) and node(s);
- b) Stop the associated out-bound data transfers of those commands that were in progress over the failing link;
- c) Issue a QUIESCE SMS to each node that was executing an affected command, providing the node remains configured (i.e., by an alternate path);
- d) When a RESPONSE SMS has been received for each QUIESCE SMS, a RESPONSE SMS is sent to the Master for the MASTER ALERT SMS when all of its affected nodes have been quiesced;

- e) Issues a QUERY NODE SMS to each node that was quiesced over each path that is intended to be used. This rebuilds the Configurator table;
- f) If the link is still configured or there is an alternative path then reissue the affected commands.

11.6.5 MASTER GENERATED ASYNC ALERT Type Handling process

The actions taken with the MASTER GENERATED ALERT TYPE field value shall depend on the ALERT CODE value as defined below.

11.6.5.1 ALL OPERATIONAL PORTS ON ALL NODES ARE IN NORMAL MODE

Any Configurator node receiving this MASTER ALERT SMS informs its upper level protocol(s) that all nodes are now capable of receiving Application frames, then invokes the General MASTER ALERT SMS Handling process (see 11.6.1).

11.6.5.2 ALL OPERATIONAL PORTS ON THE SPECIFIED NODE ARE IN NORMAL MODE

This MASTER ALERT SMS is for use in switch node applications. The receiving Configurator node does not inform its upper level protocol manager that the node identified is able to receive Application frames until all the port on the path to that node are in Normal mode, then invokes the General MASTER ALERT SMS Handling process (see 11.6.1).

11.6.5.3 RECONFIGURATION REQUIRED

This MASTER ALERT SMS follows the Total Reset of a Configurator node. Any Configurator node receiving this MASTER ALERT SMS clears its Configurator table and reconfigure the Web as described in 10.4, , then invokes the General MASTER ALERT SMS Handling process (see 11.6.1).

11.6.5.4 BOTH PORTS OF LINK ARE IN NORMAL MODE

Any Configurator node receiving this MASTER ALERT SMS does not inform its upper level protocol that the identified node is able to receive Application frames until all the ports on the path to that node are in Normal mode, then invokes the General MASTER ALERT SMS Handling process (see 11.6.1).

12 SSA message structures

The typical SSA message structure (SMS) is defined in Table 25. If the SMS CODE field value (byte 0) ≥ 80h, then the SMS is an upper level protocol SMS with the format defined in the appropriate upper level protocol. The SMS is transported in the DATA field of a frame (see 8.6).

Table 25 - Generic SSA message structure

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE							
1	(varies with SMS)							
2	TAG							
3	TAG							
4	(varies with SMS)							
...	(varies with SMS)							
31	(varies with SMS)							

An SMS has a maximum length of 32 bytes. Only the bytes defined in the appropriate SMS structure are required, and the SMS may be padded up to a maximum of 32 bytes. If the source node pads the SMS, the pad bytes shall be zero. The destination node shall ignore any pad bytes.

A path address in an SMS (e.g., RETURN PATH, ALERT PATH) shall be left-aligned in a four-byte field. The EXTEND flag function is defined in 8.5.1. The bytes of the path address following the byte with an EXTEND flag of zero are ignored.

The SMS codes that are valid for SSA-TL2 and associated frame type and node type support is defined in Table 26.

Table 26 - SSA-TL2 messages supported

SMS	SMS CODE	SMS FRAME TYPE	Node type support	
			Sent from	Received by
QUERY NODE	00h	Privileged	Configurator	all
QUERY NODE REPLY	01h	Privileged	all	Configurator
CONFIGURE PORT	02h	Privileged	Master	all
RESPONSE	03h	Privileged	all	Configurator
ASYNC ALERT	04h	Privileged	all	Master
MASTER ALERT	05h	Privileged	Master	Configurator
QUIESCE	06h	Privileged	Configurator	all
ASYNC REPLY	07h	Privileged	Master	all
QUERY PROTOCOL	08h	Privileged	Configurator	all
QUERY PROTOCOL REPLY	09h	Privileged	all	Configurator
QUERY PORT	0Ah	Privileged	Configurator	all
QUERY PORT REPLY	0Bh	Privileged	all	Configurator
QUERY SWITCH	0Ch	Privileged	Configurator	all
QUERY SWITCH REPLY	0Dh	Privileged	all	all
REQUEST SAT REGION	0Eh	Privileged	Configurator	Master
REQUEST SAT REGION REPLY	0Fh	Privileged	Master	Configurator
QUERY REGISTRATION	10h	Privileged	Master	all
DELETE RETURN PATH ID	11h	Privileged	Master	all
reserved for Privileged frames	12h-1Fh	Privileged	reserved	reserved
DATA READY	20h	Application	all	all
DATA REPLY	21h	Application	all	all
DATA REQUEST	22h	Application	all	all
DATA REDIRECT	23h	Application	all	all
reserved for Application frames	24h-7Fh	Application	reserved	reserved
Defined by ULP	see 12.2.7	defined in ULP		

See 8.1.1 for information on encoding the frame type.

12.1 SMS validation

Each SMS undergoes three levels of validation.

12.1.1 SMS code validation

The code value of an SMS shall be validated as follows in order:

- If the SMS CODE value is greater than or equal to 80h and the specified ULP is supported, then the SMS is sent to the specified ULP. The FRAME TYPE is not checked.
- If the SMS CODE value is greater than or equal to 80h and the specified ULP is not supported, then an Asynchronous Alert process shall be invoked with an ALERT CODE value of UNSUPPORTED ULP.
- If the SMS CODE value is not valid, then an Asynchronous Alert process shall be invoked with an ALERT CODE value of UNKNOWN SMS.
- If the SMS CODE value is a non-reserved value shown in Table 26, but the FRAME TYPE or NODE TYPE are incorrect, then an Asynchronous Alert process shall be invoked with an ALERT CODE value of UNKNOWN SMS.
- If the SMS CODE value corresponds to the CONFIGURE PORT, MASTER ALERT, QUERY PORT, QUERY PROTOCOL, or QUIESCE SMS, and one of these SMSs is already outstanding in the node from this Configurator node, then generate a RESPONSE SMS with a RETURN CODE field value of OVERLAPPED SMSS ATTEMPTED.
- If the SMS is a QUERY NODE, and the node has no room to store the SMS, the node shall discard the QUERY NODE SMS. The node shall not generate any type of response. A Configurator node shall employ a time-out to retry the QUERY NODE if no QUERY NODE REPLY is received.
- If Table 26 matches the SMS CODE value (excluding reserved values), the FRAME TYPE of the SMS, and the node type, then continue processing the SMS as defined in the clause 12.

12.1.2 SMS length validation

The length of an SMS shall be validated. If the number of bytes received is less than specified for the SMS the port shall invoke an Asynchronous Alert process with an ALERT CODE value of SMS TOO SHORT.

12.1.3 SMS field validation

The fields of an SMS shall be validated as described below. If the Asynchronous Alert process is invoked then the message is terminated, without any other response. The RETURN PATH field of Privileged SSA-TL2 SMSs need not be registered in the Configurator table. The SMS field validation process is as follows in order:

- a) If any reserved field is non-zero or a valid field contains a reserved value, then generate an appropriate response as defined in clause 11.2.7. This checking is not performed against the QUERY NODE and QUERY NODE REPLY SMSs.
- b) If the TAG field is invalid, or valid with an improper response, then the Configurator node logs the error and ignores the SMS. If the SMS was a data transfer SMS, the Configurator node aborts the I/O process using the ULP if the ULP is known.

12.2 Privileged SMS

12.2.1 ASYNC ALERT SMS

The ASYNC ALERT SMS defined in Table 27 is sent by a node to inform the Master of an Asynchronous Alert. The Master forwards the ASYNC ALERT SMS to each Configurator node as a MASTER ALERT SMS (see 12.2.4). The Master returns an ASYNC REPLY SMS when recovery has been completed (see 12.2.2).

Table 27 - ASYNC ALERT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (04h)							
1	PORT							
2	TAG							
3	TAG							
4	ALERT CODE							
5	ALERT CODE							
6	ALERT CODE							
7	reserved							
8	CONTROL							
9	CHANNEL							
10	CHANNEL							
11	FRAME DATA							
	FRAME DATA							
n	FRAME DATA							

The PORT field contains an unsigned integer to identify the affected port.

The TAG field contains the tag specified by the most-recent CONFIGURE PORT SMS for the affected port. It allows the Master to determine the affected port on the node that sent the Asynchronous Alert.

The ALERT CODE field consists of three parts as shown in Table 28.

Table 28 - Format of the ALERT CODE field

Byte within SMS	Byte Name	Description
4 ¹	TYPE	The general type of Asynchronous Alert generated
5 ¹	SUBTYPE	Additional information that is type dependent.
6 ¹	TYPE INFORMATION	Additional information type and subtype dependent.
NOTE ¹ - The TYPE, SUBTYPE, and TYPE INFORMATION fields in the MASTER ALERT SMS are in bytes 16, 17, and 18 respectively.		

The CONTROL field contains the control byte of the frame causing the ASYNC ALERT SMS.

The CHANNEL field contains the channel of the frame causing the ASYNC ALERT SMS.

The FRAME DATA field contains up to the first 21 bytes of the data field of the frame causing the Asynchronous Alert.

The CONTROL, CHANNEL, and FRAME DATA fields are only valid for meaning codes of UNKNOWN SMS, INVALID SMS, or PROTOCOL ERROR.

The ADDRESS field in an ASYNC ALERT SMS is copied from the RETURN PATH field in the most recent CONFIGURE PORT SMS for the affected port. Any padding bytes shall be discarded.

The acknowledgment for the ASYNC ALERT SMS is the receipt of an ASYNC REPLY SMS with the same tag value. The RETURN CODE field indicates if the ASYNC ALERT SMS was successfully parsed (see 12.2.2).

The Asynchronous Alert process shall be invoked for a Link ERP process exit if the ASYNC ALERT SMS is sent via a different port than the one that had the Link ERP process exit. If the ASYNC ALERT SMS is to be sent via the same port then the alert is discarded. The remote port is still be able to report the alert since it is nearer the Master. All other types of Asynchronous Alert are always reported.

A particular alert remains pending in the affected port until the node receives an ASYNC REPLY SMS for the corresponding ASYNC ALERT SMS. The ASYNC REPLY SMS indicates that the alert has been processed by the Master. If the node receives no ASYNC REPLY within 500 ms of sending an ASYNC ALERT SMS then it retransmits the ASYNC ALERT SMS. This cycle repeats indefinitely until an ASYNC REPLY SMS, a Total Reset, or an Absolute Reset is received. The Master shall be prepared to receive more than one ASYNC ALERT SMS for the same event and send an ASYNC REPLY SMS for each SMS. Also the affected node shall be prepared to discard an ASYNC REPLY SMS if the corresponding port has no alert pending.

The CONFIGURE PORT SMS allocates a single tag to be used by ASYNC ALERT SMS. Therefore, each port shall have only one ASYNC ALERT SMS that is awaiting an ASYNC REPLY SMS at any given time, requiring queuing of Asynchronous Alerts be supported in the port.

The generation of ASYNC ALERT SMSs is controlled by the De-queue Asynchronous Alert process (see 11.4).

The meaning code defined in Table 29 is the description assigned the combination of the TYPE, SUBTYPE, and TYPE INFO fields.

Table 29 - ALERT CODE values

TYPE	SUBTYPE	TYPE INFORMATION	ALERT CODE VALUE	Port mode
Type 00h, PORT NOW OPERATIONAL				
00h	01h	00h	PORT NOW OPERATIONAL	Privileged
Type 01h, ADDRESS ERROR				
01h	01h	00h	ADDRESSED PORT NOT OPERATIONAL	Privileged
Type 02h, UNKNOWN SMS				
02h	01h	00h	UNKNOWN SMS CODE	Unchanged
02h	02h	00h	reserved	
02h	03h	00h	SMS TOO SHORT	Unchanged
02h	04h	00h	UNSUPPORTED ULP	Unchanged
Type 03h, INVALID SMS				
03h	01h	00h	UNKNOWN RETURN PATH OR UNKNOWN RETURN PATH ID	Unchanged
03h	02h	00h	INVALID FIELD	Unchanged
Type 04h, PROTOCOL ERROR				
04h	01h	00h	SMS UNEXPECTED	Unchanged
04h	02h	00h	reserved	
04h	03h	00h	TOO MANY SMS OUTSTANDING	Unchanged
04h	04h	00h	TOO MUCH DATA RECEIVED	Unchanged
Type 05h, ENVIRONMENTAL ERROR				
05h	01h	00h	POWER FAULT	Unchanged
05h	02h	00h	CRITICAL FAN FAILURE	Unchanged
Type 06h, WARNING				
06h	01h	00h	REDUNDANT FAN FAILURE	Unchanged
06h	02h	vendor specific	FAILURE PREDICTION THRESHOLD EXCEEDED	Unchanged
06h	03h	00h	PORT NOT OPERATIONAL DUE TO HARDWARE FAULT	Wrap
06h	04h	MASTER PRIORITY field	MASTER PRIORITY CHANGED	Unchanged
06h	05h	CURRENT SPEED field	PORT OPERATING AT SLOWER THAN OPTIMAL SPEED	Unchanged
06h	06h	ULP CODE	SUPPORTED ULP LIST CHANGED	Unchanged
Type 10h to 19h, Link ERP process exits. Local = Link Status Byte sent by the node originating the ASYNC ALERT. Remote = Link Status Byte received by the node originating the ASYNC ALERT.				
10h	00h	00h	PERMANENT LINE FAULT	Privileged
11h	00h	00h	NO CHARACTERS RECEIVED	Privileged
12h	00h	00h	REMOTE PORT DISABLED	Privileged
13h	Local	00h	LINK RESET FAILED	Privileged
14h	Local	Remote	RETRY LIMIT EXCEEDED	Privileged
15h	Local	Remote	HARDWARE ERROR	Privileged
16h	Local	Remote	FRAME REJECT	Privileged
17h	Local	Remote	INVALID RETRY STATUS	Privileged
18h	Local	Remote	TIME-OUT WAITING FOR DISABLED STATE	Privileged
19h	Local	Remote	TIME-OUT WAITING FOR READY STATE	Privileged
Type 1Ah, ERROR THRESHOLDS.				
1Ah	00h	00h	ALARM THRESHOLD EXCEEDED	Unchanged
Type BCh to C0h, Master generated alerts (not used in ASYNC ALERT).				
BCh	00h	00h	ALL OPERATIONAL PORTS ON ALL NODES IN THE WEB ARE IN NORMAL MODE.	Normal
BDh	00h	00h	ALL OPERATIONAL PORTS ON THE SPECIFIED NODE ARE IN NORMAL MODE.	Normal
BEh	00h	00h	RECONFIGURATION REQUIRED	N/A
BFh	00h	00h	BOTH PORTS OF LINK ARE IN NORMAL MODE.	Normal
C0h	00h	00h	SAT REGION MEMBERSHIP ESTABLISHED	Unchanged
C0h	01h	00h	SAT REGION MEMBERSHIP DELETED	Unchanged

12.2.2 ASYNC REPLY SMS

The ASYNC REPLY SMS defined in Table 30 is sent by the Master to acknowledge an ASYNC ALERT SMS.

Table 30 - ASYNC REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (07h)							
1	RETURN CODE							
2	TAG							
3	TAG							
4	ALERT CODE							
5	ALERT CODE							
6	ALERT CODE							

The RETURN CODE field indicates the result of the original ASYNC ALERT SMS. The return code values are defined in Table 31.

Table 31 - RETURN CODE field values

Value	Description
00h	ASYNC ALERT SMS SUCCESSFULLY RECEIVED
FFh	INVALID FIELD

The TAG field identifies the SMS that is being acknowledged and is copied from the original ASYNC ALERT SMS. The tag value is not released and shall be reused for the next ASYNC ALERT SMS unless changed by a subsequent CONFIGURE PORT SMS.

The ALERT CODE field consists of three parts as shown in Table 29.

The node shall ignore any ASYNC REPLY SMS with TYPE, SUBTYPE and TYPE INFORMATION fields that do not match those of the outstanding ASYNC ALERT SMS.

If the TAG and ALERT CODE fields do not match the top entry in the Queued Asynchronous Alert table, the ASYNC REPLY SMS is ignored. If the TAG and ALERT CODE fields match the top entry in the Queued Asynchronous Alert table, then that entry is removed and the De-queue Asynchronous Alert process is invoked.

12.2.3 CONFIGURE PORT SMS

The CONFIGURE PORT SMS defined in Table 32 is sent by the Master to each port on every other node in the Web during the configuration process. The Master also issues a CONFIGURE PORT SMS to return a port to Normal mode after it has processed an Asynchronous Alert that placed the port in the Privileged mode. In both cases the destination node returns a RESPONSE SMS.

Table 32 - CONFIGURE PORT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (02h)							
1	PORT							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	reserved							
9	A QUOTA							
10	B QUOTA							
11	EUDC	REFLECT	MODE		reserved		EEDO	RACK
12	reserved							
13	ALARM THRESHOLD							
14	ALARM THRESHOLD							
14	NEGOTIATE SPEED							
15	NEGOTIATE SPEED							

A CONFIGURE PORT SMS specifies the return path and the tag to be used when a node sends the Master an ASYNC ALERT SMS. It also establishes the operating parameters of the port.

The PORT field contains an unsigned integer to identify the affected port.

The TAG field is returned in the RESPONSE SMS. The same tag is also used in an ASYNC ALERT SMS if the node subsequently reports an Asynchronous Alert.

The TAG value is assigned by the Master and shall be unique among the TAG values used by other ports in the Web as indicated by the Asynchronous Alert table (see Table 16). The receipt of the resulting RESPONSE SMS shall release the TAG value for use by SMSs other than the CONFIGURE PORT SMS, and the tag value shall no longer be considered active.

The RETURN PATH field specifies the path component that shall be placed in the address field of the resulting RESPONSE SMS and any subsequent Asynchronous Alert.

The A QUOTA field specifies the frame quota that a node may originate before it is satisfied. The value of the A QUOTA field shall be greater than zero.

The B QUOTA field specifies the maximum frame quota that a node may originate for each rotation of the SAT token. The value of the B QUOTA field shall be greater than or equal to the value of the A QUOTA field. If the B QUOTA value is less than the A QUOTA value, then the CONFIGURE PORT SMS shall be ignored and a RESPONSE SMS generated to the Configurator node issuing the CONFIGURE PORT SMS with the appropriate TAG value and a RETURN CODE of INVALID FIELD.

The EUDC bit (Enable User Defined Characters) specifies how the port handles User Defined characters. If the EUDC bit is cleared, the port shall not transmit any User Defined characters and shall discard received User Defined characters. If the EUDC bit is set, the port may transmit User Defined characters, and shall forward User Defined characters.

The REFLECT bit is one of the controls affecting the REFLECTION FLAG node function (see 10.2.5). If the node is a dual port node and the REFLECT bit is set, the REFLECTION FLAG is set for the node. If the node is a dual port node with the OPERATIONAL FLAG set for both ports and the REFLECT bit is cleared, the REFLECTION FLAG is cleared for the node. The REFLECT bit shall be ignored for single port nodes or switch nodes.

The MODE field selects the port mode as defined in Table 33.

Table 33 - MODE field values

Value	Description
00b	NO CHANGE TO THE CURRENT MODE.
01b	SET WRAP MODE.
10b	SET NORMAL MODE.
11b	SET PRIVILEGED MODE.

If a node receives a CONFIGURE PORT SMS with MODE field value of NO CHANGE TO THE CURRENT MODE and the specified port has an Asynchronous Alert pending then it shall send all pending ASYNC ALERT to the Master before sending the RESPONSE SMS for the CONFIGURE PORT. A MODE field value of SET WRAP MODE sets the Wrap mode and is only valid if the node received the CONFIGURE PORT SMS via a different port than that being configured. Otherwise the node returns a RESPONSE SMS with RETURN CODE VALUE of INVALID FIELD and the port mode is unchanged. When a port enters the Wrap mode, the port also becomes not Operational (see 10.1.4.3). Once a port has been placed in Wrap mode, the Master causes the node to exit the Wrap mode on the port by sending a CONFIGURE PORT SMS via a different port with a mode field value of SET PRIVILEGED MODE. The port shall exit the Wrap mode and invoke the Beginning Communication process (see 9.4). If a port currently configured in Wrap mode receives a CONFIGURE PORT SMS via a different port with a mode field value of SET NORMAL MODE, the node returns with a RESPONSE SMS with RETURN CODE VALUE of INVALID FIELD and the port remains in the Wrap mode.

The EEDO bit enables the Extended Distance Option. If the EDOE bit is set and the port does not support the Extended Distance Option, a RESPONSE SMS is generated with a RETURN CODE value of INVALID FIELD. If the Extended Distance Option is supported and the EDOE bit is set the ACK time-out is increased to between 100 μ s and 200 μ s, and the port may now send RR character pairs for each buffer it has available to receive frames. If the EDOE bit is cleared, the ACK time-out is set to between 25 μ s and 50 μ s, and the Extended Distance Option is disabled.

The RACK bit reduces the ACK time-out to between 25 μ s and 50 μ s. If the RACK bit is set, the ACK time-out is set to between 25 μ s and 50 μ s. If the RACK bit is cleared, no action is taken.

The ALARM THRESHOLD field specifies the number of times the Link ERP process is invoked before the Asynchronous Alert process is invoked with an ALERT CODE value of ALARM THRESHOLD EXCEEDED. An ALARM THRESHOLD field set to zero disables the alarm threshold function. When the LINK ERP COUNTER = ALARM THRESHOLD value the node shall issue an Asynchronous Alert. Subsequent increments to the LINK ERP COUNTER above the ALARM THRESHOLD value shall not invoke the Asynchronous Alert process. If a valid QUERY PORT SMS is received with the CLE flag set, the specified port shall set the LINK ERP COUNTER to zero and thereby restart the alarm threshold function (see 12.2.9).

The NEGOTIATE SPEED field defined in Table 30 indicates what speeds the port should use when speed negotiating during the Beginning Communication process (see 9.3). The changes to speed negotiation caused by the state of the NEG40 and NEG20 flags is defined in Table 31. If a port is instructed to operate at only 20 MB/sec or to operate at only 40 MB/sec, then the node behaves as a single speed device (see 9.4.1).

Table 30 - NEGOTIATE SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0	
14	reserved								
15	reserved						NEG40	NEG20	

Table 31 - NEGOTIATE SPEED field effects on speed negotiation

NEG40	NEG20	Speeds to be negotiated
Cleared	Cleared	No change to speed negotiation
Cleared	Set	Speed negotiate only at 20 MB/sec
Set	Cleared	Speed negotiate only at 40 MB/sec
Set	Set	Speed negotiate at both 40 MB/sec and at 20 MB/sec

A port only saves the parameters of the most-recent CONFIGURE PORT that it has received.

The acknowledgment for the CONFIGURE PORT SMS is the receipt of a RESPONSE SMS with the same TAG value as was used for the CONFIGURE PORT SMS. The RETURN CODE field of the RESPONSE SMS shall indicate whether or not the CONFIGURE PORT was processed successfully (see 12.2.16).

Following the generation of a RESPONSE SMS for the CONFIGURE PORT SMS, a port that is running at 20 M/sec but is capable of running at both 40 MB/sec and 20 MB/sec shall invoke the Asynchronous alert process with an ALERT CODE value of PORT OPERATING AT SLOWER THAN OPTIMAL SPEED.

12.2.4 DELETE RETURN PATH ID SMS

The DELETE RETURN PATH ID SMS defined in Table 34 is used by the Master to request the target to delete the Configurator table entry identified by the port number and Return Path ID indicated in the SMS.

Table 34 - DELETE RETURN PATH ID SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (11h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
4	RETURN PATH ID							
5	RETURN PATH ID							
6	RETURN PATH ID							
7	RETURN PATH ID							
16	reserved	PORT-1						

The TAG field is returned in the RESPONSE SMS. The TAG value is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specified the Path component that shall be placed in the Path component of the ADDRESS field of the resulting RESPONSE SMS.

The RETURN PATH ID field is the RETURN PATH ID field value sent by the target node to the Master as a result of the SEND REGISTRATION SMS that was identified by the Master as Invalid.

The PORT-1 field is the PORT-1 field value sent by the target node to the Master as a result of the SEND REGISTRATION SMS that was identified by the Master as Invalid.

The destination node shall find the Configurator table entry that meets the following conditions:

- a) matches the Return Path ID identified by the RETURN PATH ID field value;
- b) it is assigned to the port identified by the PORT-1 field value,
- c) it has a REPORTED status (see 9.2.4).

If a matching entry is found, then the STATUS field of the entry is changed to FREE and a RESPONSE SMS is sent with a RETURN CODE field value of REQUESTED FUNCTION WAS COMPLETED SUCCESSFULLY. If a matching entry is not found, then a RESPONSE SMS is generated with a RETURN CODE field value of REQUESTED FUNCTION FAILED. In either case, the TAG field value is used in the RESPONSE SMS and the RESPONSE SMS is sent to the address indicated by the RETURN PATH field value and the port the DELETE RETURN PATH ID SMS was received on.

12.2.5 MASTER ALERT SMS

The MASTER ALERT SMS defined in Table 35 is sent from the Master to each other Configurator node for the following three reasons:

- a) To forward an ASYNC ALERT SMS or an equivalent alert detected within the Master node itself. If the alert is for a Link ERP process exit then the Configurator node quiesces any I/O processes that were using the affected port before returning a RESPONSE SMS to the Master. The Configurator node invalidates the entries in the Configuration table for the affected nodes (see 12.2.1).
- b) To indicate that both ports of a link have been put into Normal mode (e.g., at power-on or following error recovery). The Configurator node returns a RESPONSE SMS to the Master immediately.
- c) To inform other Configurator nodes that they need to invoke the Configuration process to reconfigure the Web.
- d) To inform Configurators that they have been included in a SAT Region as a result of a REQUEST SAT REGION SMS from another Configurator.

Table 35 - MASTER ALERT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (05h)							
1	PORT							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	ALERT CODE							
17	ALERT CODE							
18	ALERT CODE							
19	reserved							
20	CONTROL							
21	CHANNEL							
22	CHANNEL							
23	FRAME DATA							
...	FRAME DATA							
31	FRAME DATA							

The PORT field identifies the port that invoked the Asynchronous Alert.

The TAG field is assigned by the Master and returned in the response from the destination Configurator node. It shall be unique among all the tags currently active from the Master.

The RETURN PATH field specifies the path component that shall be placed in the ADDRESS field of the RESPONSE SMS.

The UNIQUE ID field identifies the node that invoked the Asynchronous Alert.

The ALERT CODE field consists of three parts as shown in Table 28, and is copied from the same field of the associated ASYNC ALERT SMS.

The CONTROL field is copied from the same field in the associated ASYNC ALERT SMS.

The CHANNEL field is copied from the same field in the associated ASYNC ALERT SMS.

The FRAME DATA field is copied from the first nine bytes of the same field in the associated ASYNC ALERT SMS.

The acknowledgment for the MASTER ALERT SMS is the receipt of a RESPONSE SMS with the same TAG value as was used for the MASTER ALERT SMS. The return CODE field of the RESPONSE SMS shall indicate whether or not the MASTER ALERT SMS was processed successfully (see 12.2.16).

12.2.6 QUERY NODE SMS

This SMS is sent from a Configurator node to every other Operational node during the Configuration process. QUERY NODE is also used as a remote wrap test to verify the integrity of the Path.

The QUERY NODE SMS is defined in Table 36.

Table 36 - QUERY NODE SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (00h)							
1	SSA-TL VERSION (10h)							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	DR	MA	MASTER PRIORITY			reserved		

The SSA-TL VERSION field defined in Table 37 identifies the version of SSA-TL being used by the sender. If multiple SSA-TL versions are supported, the node shall report the highest level supported.

Table 37 - SSA-TL VERSION field values

Version	Description
00h	Implementations prior to standardization
01h	SSA-IA/95PH
02h	SSA-TL1 standard
03h	reserved
04h	reserved
05h	SSA-IA/95PH Plus
06h	SSA-TL1 Plus
07h-0Fh	reserved
10h	This (SSA-TL2) standard
11h-FFh	reserved

NOTE 13 - The SSA-IA/95PH, SSA-IA/95PH Plus, and SSA-TL1 Plus documents may be found on the T10 home page (<http://www.x3.org/x3t10>).

The TAG field is returned in the QUERY NODE REPLY SMS. The TAG is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the path component that shall be placed in the Configurator table entry created in response to this QUERY NODE SMS, if an entry is created. This value is used for the address field of the resulting QUERY NODE REPLY SMS, and is used for the address field of any future application SMS that utilizes this Configurator table entry.

The UNIQUE ID field contains the Unique ID of the Configurator node that issued QUERY NODE SMS.

The DR bit (Disable Registration) controls the updating of the Configurator table in the node. If the DR bit is cleared then the node shall enter the specified RETURN PATH and UNIQUE ID into its Configurator table. If a Configurator node intends to use several alternative paths to the same node then it shall issue QUERY NODE SMS with the DR bit is cleared once over each path. If the DR bit is set the Configurator table shall not be updated.

The MA bit (Master Alive) is set when the Master invokes the Master Alive process (see 13.2), and is cleared all other times. When a node receives a QUERY NODE SMS with the MA bit set, it restarts the 10 second Master Alive timer (see 13.2).

The MASTER PRIORITY field is only valid for Master capable nodes, and contains the Master Priority of the node that issues the QUERY NODE SMS. The MASTER PRIORITY field value is used by the node that receives the QUERY NODE SMS to decide if it needs to resign mastership.

12.2.7 QUERY NODE REPLY SMS

The QUERY NODE REPLY SMS as defined in Table 38 is returned when a QUERY NODE SMS is received. The QUERY NODE REPLY SMS is returned on the same port that received the corresponding QUERY NODE SMS.

Table 38 - QUERY NODE REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (01h)							
1	PORT							
2	TAG							
3	TAG							
4	UPPER LEVEL PROTOCOL							
5	ITF	MASTER PRIORITY			reserved			
6	TOTAL OTHER PORTS							
7	SSA-TL VERSION (10h)							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	RETURN PATH ID							
17	RETURN PATH ID							
18	RETURN PATH ID							
19	RETURN PATH ID							
20	P10	P20	LONG	CM	reserved			

The PORT field indicates the number of the port currently being used.

The TAG field is copied from the QUERY NODE SMS. It identifies the QUERY NODE SMS that this reply is being generated for.

The UPPER LEVEL PROTOCOL field identifies the upper-level protocol that the node shall respond to. The UPPER LEVEL PROTOCOL field shall contain a value from Table 39.

Table 39 - Upper-level protocol code values

ULP code	Protocol	Notes
00h	USE THE QUERY PROTOCOL SMS TO LIST UPPER LEVEL PROTOCOLS SUPPORTED	1
01h	SHALL RESPOND TO NO UPPER-LEVEL PROTOCOL	1
02h	Vendor specific	1,2
03h-7Fh	reserved	
80h	SSA-IA / 95SP	1,2
81h	reserved	
82h	SSA-S2P	1,2
83h	SSA-S3P	1,2
84h-FAh	reserved	
FBh-FFh	Vendor specific	2
NOTES 1) Valid in QUERY NODE REPLY. 2) Valid in QUERY PROTOCOL REPLY		

The Configurator table full (ITF) bit is set when there is no space left in the Configurator table to make an entry for the Configurator node that sent QUERY NODE SMS.

The MASTER PRIORITY field defines the priority of the node for becoming the Web Master. A value of 000b indicates that the node is a Responder node and is not capable of functioning as a Master. Any value greater than 000b indicates the node s priority for becoming the Master. A value of 001b is the lowest priority and 111b the highest

NOTE 14 - A Configurator node optionally fixes its MASTER PRIORITY at 4 (the default priority) or provides some mechanism outside of SSA to change its MASTER PRIORITY dynamically.

The TOTAL OTHER PORTS field contains a value that is one less than the number of ports implemented. If this value exceeds two, the QUERY SWITCH SMS is used to retrieve a port mask for the switch.

NOTE 15 - The TOTAL OTHER PORTS field value shall be 0 (single port), 1 (dual port) or an odd number greater than 1 (switch) as per clause 9.2.

The SSA-TL field defined in Table 37 identifies the version of SSA-TL being used by the sender. If multiple levels of SSA-TL are supported and any are numerically less that the SSA-TL VERSION field of the QUERY NODE SMS, then the highest value shall be reported in the QUERY NODE REPLY SMS that is equal to or less than the SSA-TL VERSION field in the associated QUERY NODE SMS. The following are examples:

- a) An SSA-TL1/SSA-TL2 Configurator sends a QUERY NODE SMS indicating an SSA-TL version of SSA-TL2. An SSA-TL1 only node responds with a QUERY NODE REPLY SMS indicating SSA-TL1. The Configurator now knows to use SSA-TL1 to communicate with the node.
- b) An SSA-TL1/SSA-TL2 Configurator sends a QUERY NODE SMS indicating an SSA-TL version of SSA-TL2. An SSA-TL1/SSA-TL2 or SSA-TL2 only node responds with a QUERY NODE REPLY SMS indicating SSA-TL2. The Configurator now knows to use SSA-TL2 to communicate with the node.
- c) An SSA-TL1 only Configurator sends a QUERY NODE SMS indicating an SSA-TL version of SSA-TL1. An SSA-TL1/SSA-TL2 node responds with a QUERY NODE REPLY SMS indicating SSA-TL1. The Configurator now knows to use SSA-TL1 to communicate with the node.
- d) An SSA-TL1 only Configurator sends a QUERY NODE SMS indicating an SSA-TL version of SSA-TL1. An SSA-TL2 only node responds with a QUERY NODE REPLY SMS indicating SSA-TL2. The Configurator now knows it cannot communicate with the node, and does not register with it.

The UNIQUE ID field contains the node s Unique ID.

The RETURN PATH ID contains a value created by the node and returned to the Configurator node. This field shall be used by the Configurator node in any future application SMSs that are utilizing the same path as used by the QUERY NODE SMS that caused this QUERY NODE REPLY SMS to be generated.

The P10 bit is only valid for single port and dual port nodes, Switch nodes shall clear the P10 bit. For single port or dual port nodes, the P10 bit is set if port 1 is operational, and is cleared if port 1 is not operational.

The P20 bit is only valid for single port and dual port nodes, Switch nodes shall clear the P20 bit. For single port or dual port nodes, the P20 bit is set if port 2 is operational, and is cleared if port 2 is not operational.

The LONG bit is set if any port in the node is equipped with the Extended Distance Option, and cleared otherwise. If the LONG bit is set, a QUERY PORT SMS may be used to determine which ports support the Extended Distance Option.

If the CM bit is cleared, the node that issued the QUERY NODE REPLY SMS is not currently the Master. If the CM bit is set, the node that issued the QUERY NODE REPLY SMS is currently the Master. If the current Master receives a QUERY NODE SMS with a MASTER PRIORITY field value greater than its own, then it responds with a QUERY NODE REPLY SMS with the CM bit set and its master priority in the MASTER PRIORITY field, and then resigns mastership.

The Path component of the ADDRESS field in a QUERY NODE REPLY SMS frame is copied from the RETURN PATH field in the corresponding QUERY NODE SMS. All padding bytes shall be discarded.

12.2.8 QUERY PORT SMS

The QUERY PORT SMS defined in Table 40 is sent by a Configurator node to request the error statistics of the specified port. The destination node normally returns a QUERY PORT REPLY SMS. However, an invalid field causes a RESPONSE SMS to be generated with a RETURN CODE of INVALID FIELD.

Table 40 - QUERY PORT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Ah)							
1	PORT							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	reserved							CLE

The PORT field contains the port of the link status report.

The TAG field is returned in the QUERY PORT REPLY SMS. The TAG is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the path component that shall be placed in the path component of the ADDRESS field of the resulting QUERY PORT REPLY SMS.

The CLE bit specifies whether the node shall clear the LINK ERP COUNTER. If this CLE bit is set, the port shall reset the counter to zero after sending the QUERY PORT REPLY SMS. If CLE bit is cleared, the port shall send the current value of the LINK ERP COUNTER without clearing it.

12.2.9 QUERY PORT REPLY SMS

The QUERY PORT REPLY SMS defined in Table 41 is sent from a node to a Configurator node in response to a QUERY PORT SMS.

Table 41 - QUERY PORT REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Bh)							
1	reserved							
2	TAG							
3	TAG							
4	LINK ERP ERROR COUNT							
5	LINK ERP ERROR COUNT							
6	A QUOTA							
7	B QUOTA							
8	EUDC	REFLECT	MODE		reserved		EDOC	EDOE
9	reserved							
10	ALARM THRESHOLD							
11	ALARM THRESHOLD							
12	SUPPORTED SPEED							
13	SUPPORTED SPEED							
14	CURRENT SPEED							
15	CURRENT SPEED							

The TAG field is copied from the QUERY PORT SMS. It identifies the QUERY PORT SMS that this response is being generated for.

The LINK ERP ERROR COUNT field is the count of the number of times the Link ERP has been invoked for the selected port since the last power on, Total Reset Control frame, Absolute Reset Control frame or the last QUERY PORT MS with the CLE bit is set. A value of FFFFh is returned if the LINK ERP COUNTER is greater than or equal to 65535 (the count stops incrementing at 65535).

The A QUOTA field specifies the current frame A quota for the port set by the CONFIGURE PORT SMS for the selected port.

The B QUOTA field specifies the current frame B quota for the port set by the CONFIGURE PORT SMS for the selected port.

The EUDC bit specifies the current setting of the EUDC bit for the port set by the CONFIGURE PORT SMS for the selected port.

The REFLECT bit specifies the current setting of the REFLECT FLAG for the node.

The MODE field specifies the current port mode for the port set by the CONFIGURE PORT SMS for the selected port.

The EDOC bit indicates if a port is capable of the Extended Distance Option. The EDOC bit is set if the port is capable of the Extended Distance Option and is cleared otherwise.

The EDOE bit indicates that the port currently has the Extended Distance Option enabled (see 12.2.3). The EDOE bit is set if the Extended Distance Option is enabled, and cleared otherwise.

The ALARM THRESHOLD field is the current setting of the ALARM THRESHOLD set by the CONFIGURE PORT SMS for the selected port. See 11.2.3.

The SUPPORTED SPEED field is defined in Table 42. The SUP40 bit shall be set if the port is capable of operating at 40 MB/s and cleared otherwise. The SUP20 bit shall be set if the port is capable of operating at 20 MB/s and cleared otherwise.

Table 42 - SUPPORTED SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
12	reserved							
13	reserved						SUP40	SUP20

The CURRENT SPEED field is defined in Table 43. If the port is currently operating at 40 MB/s, then the CUR40 bit shall be set and the CUR20 bit shall be cleared. If the port is currently operating at 20 MB/s, then the CUR40 bit shall be cleared and the CUR20 bit shall be set. If the port's OPERATIONAL FLAG is cleared, then both the CUR40 and the CUR20 bits shall be cleared.

Table 43 - CURRENT SPEED field

Byte	Bit 7	6	5	4	3	2	1	Bit 0
14	reserved							
15	reserved						CUR40	CUR20

12.2.10 QUERY PROTOCOL SMS

The QUERY PROTOCOL SMS defined in Table 44 is sent by a Configurator node to determine the upper-level protocols the destination node responds to. The destination node normally returns a QUERY PROTOCOL REPLY SMS. However, an invalid field causes a RESPONSE SMS to be generated with a RETURN CODE of INVALID FIELD.

Table 44 - QUERY PROTOCOL SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (08h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							

The TAG field is returned in the QUERY PROTOCOL REPLY SMS. The TAG is assigned by the Configurator node and shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the path component that shall be placed in the path component of the ADDRESS field of the resulting QUERY PROTOCOL REPLY SMS frame.

12.2.11 QUERY PROTOCOL REPLY SMS

The QUERY PROTOCOL REPLY SMS defined in Table 45 shall be returned a QUERY PROTOCOL SMS is received. The QUERY PROTOCOL REPLY SMS shall be returned via the same port that received the corresponding QUERY PROTOCOL. It indicates the particular upper-level protocols to that the node shall respond to.

Table 45 - QUERY PROTOCOL REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (09h)							
1	reserved							
2	TAG							
3	TAG							
4	PROTOCOL LIST							
	PROTOCOL LIST							
15	PROTOCOL LIST							

The TAG field is copied from the QUERY PROTOCOL SMS. It identifies the QUERY PROTOCOL SMS that this reply is being generated for.

The PROTOCOL LIST field contains one entry for each upper-level protocol to that the node shall respond to. Upper-level protocols are identified by ULP code values as defined in Table 39. Only values returned by QUERY PROTOCOL REPLY shall be used. The protocol list field is always 11 bytes in length and is zero filled and left justified (first byte is in byte 4 of the QUERY PROTOCOL REPLY SMS).

The entry for a standard protocol is a one-byte ULP code value.

The entry for a vendor-unique protocol consists of a one byte ULP code value followed by a 4-byte protocol identifier. The ULP code value for each protocol shall be unique. The first 3 bytes of the protocol identifier identify the original vendor of the protocol in the same way as in a Unique ID. The remaining byte is assigned by the original vendor to the particular protocol.

NOTE 16 - a node that communicates SCSI-2 and two vendor-unique protocols returns the following 11 bytes in Protocol list:

```

82h          /* SSA-S2P */
FBh 444543h 01h /* vendor A protocol */
FCh 49424Dh 03h /* vendor B protocol */
    
```

The protocol identifier is based on the OUI (see 10.2.1). The first three bytes are an OUI, the last byte is administered by the OUI assignee. For example, ACDE.4801h is the first protocol identified by the AC-DE-48 OUI assignee.

12.2.12 QUERY REGISTRATION SMS

The QUERY REGISTRATION SMS defined in Table 46 is used by the Master to set up a target node to transfer its current Configurator table information back to the Master, and to request the target to provide the Master with the number of bytes that will be transferred.

Table 46 - QUERY REGISTRATION SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (10h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							

The TAG field is returned in the QUERY REGISTRATION REPLY SMS. The TAG is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the path component that shall be placed in the path component of the ADDRESS field of the resulting QUERY REGISTRATION REPLY SMS and DATA READY SMS.

When a node receives a QUERY REGISTRATION SMS, it prepares to send its Configurator table data to the Configurator requesting it. The node sends a DATA READY SMS indicating the total table data size in the BYTE COUNT field. The node shall set the BYTE OFFSET field to zero. In response to the DATA READY SMS, the node requesting the table data sends a DATA REPLY SMS to begin the data transfer. The Configurator Table Data is sent as a set of records, one for each table entry. Each record contains 18 bytes, so that a data frame may contain up to 8 table entries. The format for a record and a data frame is shown in Table 47.

Table 47 - Configurator table returned data entry format

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	RETURN PATH ID							
1	RETURN PATH ID							
2	RETURN PATH ID							
3	RETURN PATH ID							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	PSTATUS	PORT-1						
17	reserved							

The PSTATUS bit is cleared if the table entry was valid, but has since been quiesced and is reserved for the Configurator identified by the UNIQUE ID field value. The PSTATUS bit is set if the table entry is currently valid, and the RETURN PATH ID field value is assigned to the Configurator node identified by the UNIQUE ID FIELD value and may be addressed by the RETURN PATH field value on the port identified by one more than the PORT-1 field value.

12.2.13 QUERY SWITCH SMS

The QUERY SWITCH SMS is sent from a Configurator node to each switch operational node during the configuration process to determine the ports that are operational.

The QUERY SWITCH SMS is defined in Table 48.

Table 48 - QUERY SWITCH SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Ch)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							

The TAG field is returned in the QUERY SWITCH REPLY SMS. The TAG is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the Path component that shall be placed in the Path component of the ADDRESS field of the resulting QUERY SWITCH REPLY SMS

The destination node normally returns a QUERY SWITCH REPLY SMS. If a field is invalid then the destination node returns a RESPONSE SMS with a RETURN CODE field of INVALID FIELD instead of QUERY SWITCH REPLY SMS.

12.2.14 QUERY SWITCH REPLY SMS

The QUERY SWITCH REPLY SMS as defined in Table 49 is returned when a QUERY SWITCH SMS is received. The QUERY SWITCH REPLY SMS is returned on the same port that received the corresponding QUERY SWITCH SMS.

Table 49 - QUERY SWITCH REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Dh)							
1	PORT							
2	TAG							
3	TAG							
4	PORT MASK							
...	PORT MASK							
19	PORT MASK						reserved	

The PORT field indicates the number of the port currently being used.

The TAG field is copied from the QUERY SWITCH SMS. It identifies the QUERY SWITCH SMS that this reply is being generated for.

The PORT MASK field is a bit vector, where a bit that is set indicates that the port is operational. Byte 4, bit 7 corresponds to port 1, byte 19, bit 2 corresponds to port 126. Byte 19, bits 1 and 0 are reserved since a switch is limited to 126 ports.

12.2.15 QUIESCE SMS

The QUIESCE SMS defined Table 50 is sent from a Configurator node to another node for the purpose of freeing RETURN PATH IDs and the resources associated with them in the transport layer for a given Unique ID. Quiesce is used during error recovery to abort all transport layer data transport processes that were originated from the Configurator node specified by the UNIQUE ID. The node aborts the associated data transfer processes, updates the Configurator table entry, notified the associated ULP(s) of the quiesce using asynchronous event notification and returns a RESPONSE SMS. ULPs such as S2P and S3P will abort I/O processes associated with the quiesced node.

Table 50 - QUIESCE SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (06h)							
1	RCTE	reserved						
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							

When the RCTE bit (remove Configurator table entry) is set, the node shall free all the Configurator table entries associated with the UNIQUE ID field. When the RCTE bit is cleared, the node shall hold one entry and free all other entries in the Configurator table entry associated with UNIQUE ID field (see 10.2.4), guaranteeing at least one table entry will be available when the node issuing the quiesce re-registers.

NOTE 17 - It is recommended that the RCTE bit only be set by a Configurator node quiescing its own I/O processes, by the Master or by the resigning Master node.

The TAG field is assigned by the Configurator node and returned in the RESPONSE SMS from the node. The TAG is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the path component that shall be placed in the ADDRESS field of the RESPONSE SMS.

The UNIQUE ID field specifies the Unique ID of the Configurator node whose RETURN PATH ID resources are to be released.

Specifying the Configurator node with a Unique ID rather than a return path allows the Master to issue a third party Quiesce on behalf of a missing Configurator node.

The node shall search its Configurator table to convert the Unique ID to a return path before selecting the data transport processes to be aborted. If the Configurator node is using multiple paths to the same node then there is more than one return path for the given Unique ID and the node shall Quiesce the data transport processes for all paths.

If a Configurator node is using multiple paths to the same node then it shall issue a QUIESCE SMS over each path that remains connected. When it receives the RESPONSE SMS for a particular path the Configurator node safely assumes that no frames remain in the Web for I/O processes issued over that path.

In order to release Configurator table space on the other node, an ordinary QUIESCE SMS shall be issued by Configurators that have lost a path to another node, as indicated by a Master Alert. The node will then re-register over any paths it intends to continue using. A QUIESCE SMS with the RCTE bit set is sent by a node to release all of the Configurator table space on the other node when no more additional I/O to that node is needed. This policy of registering, performing I/O and un-registering in a multi-initiator system maximizes the number of nodes that may share access to a given node.

The acknowledgment for the QUIESCE SMS is the receipt of a RESPONSE SMS with the same tag value as was used for the QUIESCE SMS. The RETURN CODE field of the RESPONSE SMS shall indicate whether or not the quiesce was processed successfully (see 12.2.16).

12.2.16 REQUEST SAT REGION SMS

The REQUEST SAT REGION SMS defined in Table 51 is used to by a Configurator to request SAT region control from the Master.

Table 51 - REQUEST SAT REGION SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Eh)							
1	SSA-TL VERSION (10h)							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	ENDPOINT #1 UNIQUE ID							
...	ENDPOINT #1 UNIQUE ID							
15	ENDPOINT #1 UNIQUE ID							
16	ENDPOINT #2 UNIQUE ID							
...	ENDPOINT #2 UNIQUE ID							
23	ENDPOINT #2 UNIQUE ID							
24	reserved						SAT OPCODE	
25	REGION IDENTIFIER							
26	ENDPOINT #1 PORT NUMBER							
27	ENDPOINT #2 PORT NUMBER							
28	A QUOTA							
29	B QUOTA							

The SSA-TL VERSION field is defined in Table 37, and is the version of the transport layer currently being used between the node and the Master.

The TAG field is returned by the Master in the REQUEST SAT REGION REPLY SMS in response to the REQUEST SAT REGION SMS. The TAG value is assigned by the Configurator node and it shall be unique among the TAG values that are currently active from that Configurator node.

The RETURN PATH field specifies the PATH component that shall be placed in the ADDRESS field of the REQUEST SAT REGION REPLY SMS frame.

The ENDPOINT #1 UNIQUE ID field is the Unique ID of the node that is the first endpoint of the requested SAT region when used with an SAT OPCODE field value of CREATE SAT REGION or EXTEND SAT REGION. The ENDPOINT #1 UNIQUE ID field is ignored if the SAT OPCODE field is DELETE SAT REGION.

The ENDPOINT #2 UNIQUE ID field is the Unique ID of the node that is the second endpoint of the requested SAT region for a SAT OPCODE value of either CREATE SAT REGION or EXTEND SAT REGION. The ENDPOINT #2 UNIQUE ID field is ignored for a SAT OPCODE value of DELETE SAT REGION.

The SAT OPCODE field specified in Table 52 identifies the operation requested by the REQUEST SAT REGION SMS.

Table 52 - SAT OPCODE field

Value	Description
00b	CREATE SAT REGION
01b	EXTEND SAT REGION
10b	DELETE SAT REGION
11b	reserved

The SAT OPCODE field value of CREATE SAT REGION requests a new SAT region identified by the two specified endpoints be created and be associated with the value in the REGION IDENTIFIER field. An endpoint id identified by the appropriate endpoint unique ID and port fields.

The SAT OPCODE field value of EXTEND A REGION requests that the SAT region identified by the two specified endpoints be combined with the SAT region identified by the region identifier field.

The SAT OPCODE field value of DELETE THE REGION causes the region identified by the REGION IDENTIFIER field to be deleted. The nodes in the deleted region return to the general pool of nodes that existed prior to the creation of the SAT region.

The REGION IDENTIFIER field identifies the region being extended (if SAT OPCODE field is EXTEND SAT REGION) or deleted (if SAT OPCODE field is DELETE SAT REGION). The REGION IDENTIFIER field is ignored if the SAT OPCODE field is CREATE SAT REGION).

The ENDPOINT #1 PORT NUMBER field is a number that represents port numbers from 1 to 128 on the node indicated by the ENDPOINT #1 UNIQUE ID field for a SAT OPCODE value of CREATE SAT REGION or EXTEND SAT REGION. Since SAT regions may end on Switch ports, an 8-bit field is required. The ENDPOINT #1 PORT NUMBER field is ignored for a SAT OPCODE value of DELETE SAT REGION.

The ENDPOINT #2 PORT NUMBER field is a number that represents port numbers from 1 to 128 on the node indicated by the ENDPOINT #2 UNIQUE ID field for a SAT OPCODE value of CREATE SAT REGION or EXTEND SAT REGION. Since SAT regions may end on Switch ports, an 8-bit field is required. The ENDPOINT #2 PORT NUMBER field is ignored for a SAT OPCODE value of DELETE SAT REGION.

The A QUOTA field is the value the Master shall use in the A QUOTA field of the CONFIGURE PORT SMSs sent to the nodes that are being grouped into the SAT Region. All nodes in the region have their A QUOTA values set to this value upon the successful creation or extension of a SAT region.

The B QUOTA field is the value the Master shall use in the B QUOTA field of the CONFIGURE PORT SMSs sent to the nodes that are being grouped into the SAT Region. All nodes in the region have their B QUOTA values set to this value upon the successful creation or extension of a SAT region.

Several rules apply to SAT regions:

- a) SAT Regions do not nest. No SAT region shall overlap or be enclosed by another.
- b) If the Master detects that a web change has caused a SAT region to contain no Configutors, it deletes the SAT region and return the nodes to the pre-region configuration.
- c) When Configutors membership in a region changes due to a request they did not originate, the Master notifies them of the change using a Master Alert (see 12.2.4) indicating addition to or deletion from a region.
- d) A SAT region shall contain more than one node.

12.2.17 REQUEST SAT REGION REPLY SMS

The REQUEST SAT REGION REPLY SMS defined in Table 53 is originated by a Master in response to a REQUEST SAT REGION SMS.

Table 53 - REQUEST SAT REGION REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (0Fh)							
1	SSA-TL VERSION (10h)							
2	TAG							
3	TAG							
4	REGION IDENTIFIER							
5	RETURN CODE							

The SSA-TL VERSION field is defined in Table 37, and is the version of the transport layer currently being used between the Master and destination node.

The TAG field is copied from the REQUEST SAT REGION SMS. It identified the REQUEST SAT REGION SMS this reply is being generated for,

The REGION IDENTIFIER field is returned by the Master and is used by the Master to identify the SAT region. If the node receiving the reply was the successful creator, this serves as a verification. If the node was attempting a create and failed, the REGION IDENTIFIER field may be used in a subsequent extend command to add nodes to the region. If the node performed a successful extend operation, the REGION IDENTIFIER field contains the identifier of the original region. If the extend operation failed, the REGION IDENTIFIER field is ignored.

The RETURN CODE field indicates the result of the REQUEST SAT REGION SMS as defined in Table 54.

Table 54 - RETURN CODE values

Value	Description
00h	SUCCESSFUL OPERATION
01h	SUCCESS WITH WARNING: CREATED OR EXTENDED REGION DOES NOT INCLUDE REQUESTER.
02h	SUCCESS WITH WARNING: DELETED REGION DID NOT INCLUDE REQUESTER.
03h	FAILURE SAT REGIONS NOT SUPPORTED
04h	FAILURE REQUESTED CREATE OR EXTEND OVERLAPS EXISTING REGION
05h	FAILURE NO PATH BETWEEN ENDPOINTS FOUND.
06h	SMS REJECTED, INVALID ENDPOINT, REGION IDENTIFIER OR PORT NUMBER

12.2.18 RESPONSE SMS

The RESPONSE SMS defined in Table 55 is used to acknowledge the CONFIGURE PORT, MASTER ALERT, and QUIESCE SMSs. In addition, in the case of invalid fields, the RESPONSE SMS is used to acknowledge the QUERY PORT, QUERY PROTOCOL and DATA REPLY SMS.

Table 55 - RESPONSE SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (03h)							
1	RETURN CODE							
2	TAG							
3	TAG							

The RETURN CODE field indicates the result of the original SMS as defined in Table 56. The meaning of a return code is not dependent on the original SMS.

Table 56 - RETURN CODE values

Type	Description
00h	REQUESTED FUNCTION WAS COMPLETED SUCCESSFULLY
01h	NO I/O PROCESS FOUND
02h	UNIQUE ID NOT FOUND
03h	INVALID RETURN PATH
04h	OVERLAPPED SMS ATTEMPTED
05h-FDh	reserved
FEh	REQUESTED FUNCTION FAILED
FFh	INVALID FIELD

The INVALID FIELD RETURN CODE value indicates that the SMS was parsed successfully but one or more fields was invalid (see 12.1, for the action taken when an SMS cannot be parsed).

Each SMS definition specifies all the values of return code that are applicable to that SMS.

The TAG field is copied from the original SMS. It identifies the SMS that is being acknowledged.

The Path component of the ADDRESS field for the RESPONSE SMS frame is copied from the RETURN PATH field in the SMS being acknowledged. All padding bytes shall be discarded.

12.3 Application SMS

The FRAME TYPE field in the CONTROL field of a frame is set to indicate an Application frame (see 7.2).

The CHANNEL address in all SMS frames is 00h.

When an SMS contains a single byte Channel as a parameter it is left-aligned in a 2-byte field, and the second byte is ignored.

12.3.1 DATA READY SMS

The DATA READY SMS defined in Table 57 is sent before the outbound transfer of data. The receiving node returns one or more DATA REPLY SMSs (or DATA REDIRECT SMSs if third party transfers are required) for each DATA READY SMS. The node optionally uses more than one DATA READY SMS to transfer all of the data for a I/O process. A DATA READY SMS is considered outstanding from the time it is sent until the last Data frame associated with it is sent.

Table 57 - DATA READY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (20h)							
1	reserved							
2	TAG							
3	TAG							
4	BYTE OFFSET							
5	BYTE OFFSET							
6	BYTE OFFSET							
7	BYTE OFFSET							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	reserved							
13	reserved							
14	TRANSFER ID							
15	TRANSFER ID							

The TAG field contains a copy of the TAG field in the ULP command SMS and allows the node to associate this SMS with the correct I/O process.

The BYTE OFFSET field indicates the byte offset of the first byte to be transferred, relative to first byte requested by the I/O process. The value of the BYTE OFFSET field shall be a multiple of eight.

The BYTE COUNT field specifies the number of bytes to be transferred. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the command.

The sum of the BYTE COUNT and BYTE OFFSET fields in a DATA READY SMS shall not exceed the transfer length of the I/O process. Checking for violation of this rule is optional.

NOTE 18 - Since data maybe retransmitted, the Configurator node cannot compare the sum of the BYTE COUNT fields in all the associated DATA READY SMSs with the expected total byte count of the I/O process.

The TRANSFER ID field is a unique identifier generated by the node to identify the DATA READY SMS. The value shall be unique among all outstanding DATA READY SMSs between a given pair of nodes. The node receiving the DATA READY SMS shall copy the TRANSFER ID into all associated DATA REPLY or DATA REDIRECT SMSs to associate the DATA REPLY or DATA REDIRECT SMS to a DATA READY SMS.

If the DDRM bit is cleared, the protocol for the DATA READY and DATA REPLY (or DATA REDIRECT) SMSs is as follows in order:

- a) When the node is ready to transfer data it sends a DATA READY SMS. This specifies how many bytes it currently has available and the starting offset within the data requested by the command along with a transfer ID.
- b) If the BYTE COUNT, BYTE OFFSET, or TRANSFER ID do not conform to the rules above, then the receiving node aborts the I/O process as allowed by the upper level protocol. Otherwise the receiving node allocates a channel to receive the data.
- c) If the data is to be sent to a third party, the node that received the DATA READY SMS waits to receive a DATA REQUEST SMS from the third party node. This DATA REQUEST SMS indicates how many bytes the third party node is willing to accept, and the channel it wishes to accept them on.
- d) The node receiving the DATA READY SMS responds with a DATA REPLY or DATA REDIRECT SMS indicating the CHANNEL and BYTE COUNT that it or the third party is willing to currently accept (and OUT PATH and OUT PORT in the case of third party transfers) and copies the TRANSFER ID from the associated DATA READY SMS.
- e) The node sends Data frames addressed to the specified channel, and (in the case of a DATA REDIRECT) to the correct path.
- f) If all of the data offered by the previous DATA READY SMS has not been transferred then the node returns to step 3 using the same or a different channel. The node may send a DATA REPLY or DATA REDIRECT SMS for an outstanding DATA READY SMS whenever there is at most one outstanding DATA REPLY or DATA REDIRECT SMS for that DATA READY SMS. A DATA REPLY or DATA REDIRECT SMS shall not span multiple DATA READY SMSs.
- g) If all of the data requested by the I/O process has not been transferred then the node returns to step 1.
- h) The node may send a DATA READY SMS when either there are no outstanding DATA READY SMSs for the command, or, at least one DATA REPLY or DATA REDIRECT has been received for each DATA READY that is outstanding for the command.

A Data frame shall not span multiple DATA REPLY SMSs, DATA REDIRECT SMSs or DATA REQUEST SMSs.

12.3.2 DATA REPLY SMS

The DATA REPLY SMS defined in Table 58 is sent in response to a DATA READY SMS. The node receiving the DATA REPLY SMS sends the requested Data frames. More than one DATA REPLY SMS may be sent in response to a single DATA READY SMS. A DATA REPLY SMS is considered outstanding from the time it is sent until the last Data frame associated with it is received.

Table 58 - DATA REPLY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (21h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH ID							
5	RETURN PATH ID							
6	RETURN PATH ID							
7	RETURN PATH ID							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	CHANNEL							
13	CHANNEL							
14	TRANSFER ID							
15	TRANSFER ID							

The TAG field is copied from the associated DATA READY SMS.

The RETURN PATH ID field identifies the node that issued the DATA REPLY SMS. The RETURN PATH ID field and TAG field serve to identify the I/O process. If the RETURN PATH ID field does not match the corresponding field of the ULP command SMS, an Asynchronous Alert is generated with an ALERT CODE value of INVALID RETURN PATH OR RETURN PATH ID. The node then waits for another DATA REPLY SMS or until the I/O process is terminated. The RETURN PATH ID field specifies the node issuing the DATA REPLY SMS, since tags may be redundant across Configurator nodes.

If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The BYTE COUNT field specifies the number of bytes that the node is willing to accept during the transfer associated with this DATA REPLY SMS. If it is less than the BYTE COUNT field of the corresponding DATA READY SMS, then the node shall send another DATA REPLY SMS when it is ready to accept more data. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the command. If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The CHANNEL field specifies the channel component of the ADDRESS field for the Data frames. The path component is obtained from the RETURN PATH field.

The DATA REPLY SMSs for a single DATA READY SMS may contain different CHANNEL values. If two consecutive DATA REPLY SMSs for the same DATA READY SMS contain the same CHANNEL value, then the node shall complete the first DATA REPLY SMS's transfer before beginning the second DATA REPLY SMS's transfer. If two consecutive DATA REPLY SMSs for the same DATA READY SMS contain different CHANNEL values, the node may overlap the data transfers.

The TRANSFER ID field is copied from the TRANSFER ID field of the associated DATA READY SMS.

In the case of an invalid field other than RETURN PATH ID, then the DATA REPLY SMS is acknowledged with a RESPONSE SMS with a RETURN CODE of INVALID FIELD. In the case of an invalid RETURN PATH ID field value, an SMS UNEXPECTED, or TOO MANY SMS OUTSTANDING, an ASYNC ALERT SMS is generated with the appropriate meaning codes. The associated I/O processes identified by the tag value (if any) is not terminated, instead the DATA REPLY SMS is rejected with a RESPONSE SMS or an ASYNC ALERT SMS.

12.3.3 DATA REQUEST SMS

The DATA REQUEST SMS defined in Table 59 is sent from by a node to request the transfer of the data for the specified channel. Multiple DATA REQUEST SMSs may be used to transfer the data for an I/O process. A DATA REQUEST SMS is considered outstanding from the time it is sent until the last Data frame associated with it is received.

Table 59 - DATA REQUEST SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (22h)							
1	reserved							
2	TAG							
3	TAG							
4	BYTE OFFSET							
5	BYTE OFFSET							
6	BYTE OFFSET							
7	BYTE OFFSET							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	CHANNEL							
13	CHANNEL							

The TAG field contains a copy of the tag field in the ULP command SMS and allows the node to associate this SMS with the correct I/O process.

The BYTE OFFSET field indicates the byte offset of the first byte to be transferred, relative to first byte requested by the I/O process. The value of the BYTE OFFSET field shall be a multiple of eight.

The BYTE COUNT field specifies the number of bytes the node is currently requesting. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the ULP command SMS.

The sum of the BYTE COUNT and BYTE OFFSET field values in a DATA REQUEST SMS shall not exceed the transfer length of the I/O process. Checking of this rule is optional.

NOTE 19 - Since data may be retransmitted, the Configurator node cannot compare the sum of the BYTE COUNT fields in all the associated DATA REQUEST SMSs with the expected total byte count of the I/O process.

The CHANNEL field specifies the channel component of the address field to be inserted in the Data frames. The DATA REQUEST SMSs associated with a single I/O process may contain different CHANNEL values. If two consecutive DATA REQUEST SMSs for the same I/O process contain the same CHANNEL value, then the first DATA REQUEST SMS's transfer shall be completed before beginning the second DATA REQUEST SMS's transfer. If two consecutive DATA REQUEST SMSs for the same command contain different CHANNEL values, the data transfers may overlap.

The protocol for the DATA REQUEST SMS is as follows in order:

- a) When the node is ready to receive data it allocates a channel and sends a DATA REQUEST. This specifies the node channel address, how many bytes the node is requesting and the starting offset within the data specified by the I/O process. If the BYTE COUNT or the BYTE OFFSET does not conform to the rules above then the I/O process may be aborted as specified by the upper level protocol.
- b) The node receiving the DATA REQUEST SMS responds by sending Data frames addressed to the specified CHANNEL. Although write data does not need to follow the same path to the node used by the I/O process, it shall arrive at the same port.

- c) If all the data specified by the I/O process has not been transferred then the node returns to step 1) after there are no outstanding DATA REQUEST SMSs for the I/O process, or at least one Data frame has been received for each DATA REQUEST SMS that is outstanding for the I/O process.
- d) A Data frame shall not span multiple DATA REQUEST SMS.

12.3.4 DATA REDIRECT SMS

The DATA REDIRECT SMS defined in Table 60 is sent in response to a DATA READY SMS, when the data is to be sent to a node other than the one that issued the command. The node receiving the DATA REDIRECT SMS sends the requested Data frames from the port specified and along the path indicated in the SMS. More than one DATA REDIRECT SMS may be sent in response to a single DATA READY SMS. A DATA REDIRECT SMS is considered outstanding from the time it is sent until the last Data frame associated with it is received by the third party.

Table 60 - DATA REDIRECT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (23h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH ID							
5	RETURN PATH ID							
6	RETURN PATH ID							
7	RETURN PATH ID							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	CHANNEL							
13	CHANNEL							
14	TRANSFER ID							
15	TRANSFER ID							
16	OUT PATH							
17	OUT PATH							
18	OUT PATH							
19	OUT PATH							
20	OUT PORT							

The TAG field is copied from the associated DATA READY SMS.

The RETURN PATH ID field identifies the node that issued the DATA REDIRECT SMS. The RETURN PATH ID field and TAG field serve to identify the I/O process. If the RETURN PATH ID field does not match the corresponding field of the ULP command SMS, an Asynchronous Alert is generated with an ALERT CODE value of INVALID RETURN PATH OR RETURN PATH ID. The node then waits for another DATA REDIRECT SMS or until the I/O process is terminated. The RETURN PATH ID field specifies the node issuing the DATA REDIRECT SMS, since tags may be redundant across Configurator nodes.

If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The BYTE COUNT field specifies the number of bytes that the node is willing to accept during the transfer associated with this DATA REDIRECT SMS. If it is less than the BYTE COUNT field of the corresponding DATA READY SMS, then the node shall send another DATA REDIRECT SMS when it is ready to accept more data. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the command. If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The CHANNEL field specifies the channel component of the ADDRESS field for the Data frames. The path component is obtained from the OUT PATH field.

The TRANSFER ID field is a unique identifier that was generated by the node that sent the DATA READY SMS. The value is unique among all outstanding DATA READY SMSs between a given pair of nodes. The node receiving the DATA READY SMS shall copy the TRANSFER ID into all associated DATA REDIRECT SMSs to associate the DATA REDIRECT SMS to a DATA READY SMS.

The OUT PATH field specifies the path component of the ADDRESS field for the Data frames. The channel component is obtained from the CHANNEL field.

The OUT PORT field specifies the port that should be used to send the Data frames.

The DATA REDIRECT SMSs for a single DATA READY SMS may contain different CHANNEL values. If two consecutive DATA REDIRECT SMSs for the same DATA READY SMS contain the same CHANNEL value, then the node shall complete the first DATA REDIRECT SMS's transfer before beginning the second DATA REDIRECT SMS's transfer. If two consecutive DATA REDIRECT SMSs for the same DATA READY SMS contain different CHANNEL values, the node may overlap the data transfers.

In the case of an invalid field other than RETURN PATH ID, then the DATA REDIRECT SMS is acknowledged with a RESPONSE SMS with a RETURN CODE of INVALID FIELD. In the case of an invalid RETURN PATH ID field value, an SMS UNEXPECTED, or TOO MANY SMS OUTSTANDING, an ASYNC ALERT SMS is generated with the appropriate meaning codes. The associated I/O processes identified by the tag value (if any) is not terminated, instead the DATA REDIRECT SMS is rejected with a RESPONSE SMS or an ASYNC ALERT SMS.

13 Transport layer processes

The transport layer processes are described in the following clauses:

- a) The Power-on reset process is defined in 10.6.3;
- b) The Reset process is defined in 10.6;
- c) The Beginning Communication process is defined in 10.3;
- d) The Configuration process is defined in 10.5.1;
- e) The Asynchronous Alert process is defined in 11.3;
- f) The De-queue Asynchronous Alert process is defined in 11.4;
- g) The Master Alert Handling process is defined in 11.6;
- h) The SAT Region process is defined in 13.3.
- i) The Configurator Table Maintenance process is defined in 13.4

13.1 Master Negotiation process

In a web with a stable topology after power-up and after the configuration process, the Configurator node with the highest MASTER PRIORITY is elected to be the Master (see 12.2.6 and 12.2.7). If more than one Configurator node is set at the same highest MASTER PRIORITY, then the Configurator node with the highest Unique ID among them is elected to be the Master.

Each Configurator node compares its own MASTER PRIORITY with the MASTER PRIORITY it receives from each other Configurator node in the QUERY NODE REPLY SMS, or the QUERY NODE SMS from another Configurator, whichever arrives first. Each Configurator node in the group with the highest MASTER PRIORITY then compares its own Unique ID with those of the other Configurator nodes in the group. The Configurator node within the group with the highest Unique ID becomes the Master.

If two sub-Webs that both contain a Configurator node are joined by a new link, then the Master of each sub-web invokes the Configuration process for the new link. When each Master received the QUERY NODE SMS from the other, they compare their MASTER PRIORITY and UNIQUE ID values. The Master node with the lower MASTER PRIORITY or Unique ID becomes a Resigning Master and notifies the other Master of this change in status by setting the CM bit in the QUERY NODE REPLY SMS. The Resigning Master no longer responds to new ASYNC ALERT SMSs. The Resigning Master then completes the processing of all pending ASYNC ALERT SMSs, and then performs a Formal Resignation to the Pending Master by sending it a QUERY NODE SMS with the MA bit

set. During the resignation procedure, the Master node with the higher MASTER PRIORITY or Unique ID becomes the Pending Master of the sub web containing the Resigning Master, sends a QUERY NODE REPLY with the CM bit set to the Resigning Master and continues acting as the Master of its own sub-Web. When the Pending Master receives the Formal Resignation from the Resigning Master it becomes the Assuming Master of the Resigning Master's sub-Web and invokes the Configuration Process for the new link. A Pending Master shall wait for 5 seconds to receive a Formal Resignation. If one has not been received in 5 seconds, then the Pending Master shall send a QUERY NODE SMS with the CM bit set to the Resigning Master. If a QUERY NODE REPLY SMS is not received within 5 seconds, then the Assuming Master shall issue a Total Reset to the node, and invoke the Configuration process for the sub Web. If a QUERY NODE REPLY SMS is received with the CM bit cleared, the Pending Master treats this as a Formal Resignation. If the CM bit is set, the Pending Master re-starts the timer. Should the second attempt also fail, then the Resigning Master node is reset and configuration continues.

If a Resigning Master receives one or more additional QUERY NODE SMSs from links added after it became a Resigning Master, then it shall return QUERY NODE REPLY SMSs with the CM bit set, and perform a Formal Resignation to each that would have won Mastership. This ensures that whatever node finally becomes Web Master will wait until the Resigning Master has completed its processing of pending ASYNC ALERT SMSs prior to configuration.

If a Pending Master receives one or more QUERY NODE SMSs that would cause it to lose the Master Election, it becomes a Resigning Master. It stops acknowledging ASYNC ALERT SMSs, completes processing of all pending ASYNC ALERT SMSs, and formally resigns to all winning Masters. Any Formal Resignations received are answered with QUERY NODE REPLY SMSs with the CM bit cleared.

If a sub-web containing no Configutors is attached to a sub-web containing a Master, the Master shall invoke the Web Reset process on the newly attached Sub-web prior to notifying other Configutors of the attachment.

The Pending Master node that determines that it now has the highest MASTER PRIORITY and Unique ID of the remaining Configutor nodes as described above becomes the Pending Master. The Pending Master then issues a CONFIGURE PORT SMS to each remaining registered Configutor node to indicate that it is the Master.

The master Negotiation process is invoked in three cases as follows and in the following sub-clauses:

- a) Two sub-webs are connected resulting in a web with two Masters;
- b) A master lowers his master priority;
- c) A non-master raises his master priority

13.1.1 Two sub-webs are connected resulting in a web with two Masters

A new link may connect two sub-webs, each with its own Master, resulting in a web containing two Masters. This case may be determined by a Master walking the newly connected portion of the Web and receiving a QUERY NODE REPLY SMS with the CM bit set indicating another Master.

To avoid ambiguity, the following terms are used:

- a) Sub-Web is a string of nodes with target functions, initiator functions, or a combination of both that are added to another web as a unit.
- b) Resignation Request is received by a Master or Resigning master of a sub-web, this is a QUERY NODE SMS that contains either a higher master priority than the receiving node's, or the same Master priority and a higher Unique ID. A resignation request is answered with a QUERY NODE REPLY with the CM bit set.
- c) Pending Resignation is whenever a QUERY NODE REPLY SMS is sent with the CM bit set. The Pending Resignation is satisfied when the Formal Resignation is performed.
- d) Resigning Master is the Master node of a Sub-Web that has received one or more Resignation Requests. It no longer accepts ASYNC ALERT SMSs, but has not yet finished processing those received prior to the QUERY NODE SMS.
- e) Formal Resignation is when a Resigning Master Formally Resigns to a node with which it has completed a Resignation Request transaction at an earlier time when it sends a QUERY NODE SMS with the MA bit set to that node.

- f) Pending Master is the current Master of a sub-web that, after walking a newly attached sub web containing one or more active Masters, has discovered that it has won mastership of that new sub web. It waits to assume mastership of the new sub-web until all of the resigning masters formally resign.
- g) Assuming Master is a Pending Master that has received all of the Formal Resignations from the Resigning Masters is the Assuming Master until it has completed the web configuration and is then the web Master. The Assuming Master may receive a Resignation Request and not become the web Master.

The Master negotiation process proceeds as follows:

- a) If two sub webs containing Masters are joined, then the master of each shall walk the other sub web to determine Mastership of the new joined web. The configuration of the new web section by the winning master does not proceed until the old Master of the sub web has completed its handling of outstanding ASYNC ALERTS and formally resigned to the new Master.
- b) When the Master of a sub-web , or a pending master receives a Resignation Request, it becomes a Resigning Master. It immediately stops accepting new ASYNC ALERT SMSs, sends a QUERY NODE REPLY SMS with the CM bit set, logs a Pending Resignation, and continues finishing the actions required by the remaining pending ASYNC ALERT SMSs. Any QUERY NODE SMSs that would not require resignation are always answered with QUERY NODE REPLY SMSs with the CM bit cleared.
- c) When a resigning master has completed all actions required by pending ASYNC ALERT SMSs, it performs Formal Resignations for all logged Pending Resignations. After resignation, the node will respond to any subsequent QUERY NODE SMSs by issuing QUERY NODE REPLY SMSs with the CM bit cleared.
- d) A Pending Master will wait until all Resigning Masters have Formally Resigned prior to configuring the ports on the new Sub-Web. When it has received all Formal Resignations from all Resigning Masters, it becomes the Assuming Master. The Assuming Master then Configures the ports in the web.
- e) If the Assuming Master receives an ASYNC ALERT SMS with an ALERT CODE value of PORT NOW OPERATIONAL, from a port during configuration, it shall check its Topology table to see if there has been a new sub-web added to the web. If not, then the ASYNC ALERT was an old queued one and the Assuming Master replies to it, discards it, and continues its configuration of the ports in the web. If there has been a new sub-web added, then the Assuming Master walks the new sub-web.
- f) Whenever a sub-web containing no Configutors is added to an existing web, the master of the existing web shall reset the nodes in the sub-web. The sub-web reset is required since the sub-web had no master when attached, and there may be outstanding commands and stale initiator table entries in the nodes.
- g) If the Assuming Master receives a Resignation Request from a node in a newly attached sub-web, it sends a QUERY NODE REPLY SMS with the CM bit set, and becomes a Resigning Master. Since the processing of the ASYNC ALERT SMS caused by the addition of the sub web includes walking the new web, this will be completed prior to resignation. If configuration of ports was in progress, this operation ceases since the new master will re-configure all of the web ports.
- h) Each Pending Master shall start a timer when it receives a reply confirming the receipt of a Resignation Request. If the Resigning Master has not issued a Formal Resignation in 5 seconds, then the Pending Master shall send a QUERY NODE SMS to the Resigning Master. If the Resigning Master returns a QUERY NODE REPLY SMS with the CM bit set, the Pending Master restarts the timer. If the Resigning Master returns a QUERY NODE REPLY SMS with the CM bit cleared, then the Pending Master accepts this reply as the Formal Resignation, but still replies to any subsequent Formal Resignations from that node. If no reply is received within another 5 seconds, the Pending Master sends a Total Reset frame to the node, followed by re-configuration of the node.

13.1.2 A master lowers his master priority;

When the current Master changes its MASTER PRIORITY the following actions are taken in the order shown:

- a) If the current Master determines that it needs to resign then it finishes processing any ASYNC ALERT SMSs previously received and it does not send an ASYNC REPLY SMS to any future ASYNC ALERT SMS.

- b) The current Master forwards its new MASTER PRIORITY to each other Configurator node in a MASTER ALERT SMS.
- c) Each other Configurator node enters the new MASTER PRIORITY into its Configuration table and returns a RESPONSE SMS.
- d) If another Configurator node determines that it needs to become the new Master then it issues a CONFIGURE PORT SMS to each port in the Web. Each port returns a RESPONSE SMS. The Configurator node is then the new Master.

13.1.3 A non-master raises his master priority

When a Configurator node (other than the current Master) changes its MASTER PRIORITY, the following process is followed in order:

- a) The Configurator node sends the current Master an ASYNC ALERT SMS containing its new MASTER PRIORITY.
- b) The current Master enters the new MASTER PRIORITY into its Configuration table and forwards it to each other Configurator node in a MASTER ALERT SMS.
- c) Each other Configurator node enters the new MASTER PRIORITY into its Configuration table and returns a RESPONSE SMS.
- d) If the current Master determines that it needs to resign then it finishes processing any ASYNC ALERT SMSs previously received and it does not send an ASYNC REPLY SMS to any future ASYNC ALERT SMSs.
- e) If the Configurator node that changed its MASTER PRIORITY determines that it needs to become the new Master, then it waits to receive the ASYNC REPLY SMS before issuing a CONFIGURE PORT SMS to each port in the Web. It is then the new Master.

13.2 Master Alive process

The loss of the Master is detected by all Configurator nodes performing the Master Alive process described as follows in order:

- a) The Master sends a QUERY NODE SMS with the MA bit set to the Configurator node with the next lower Unique ID every 3 seconds. If the Master is the lowest Unique ID, it sends a QUERY NODE SMS with the MA bit set to the Configurator node with the highest Unique ID.
- b) When a Configurator node, that is not the Master, receives a QUERY NODE SMS with the MA bit set, it sends a QUERY NODE SMS with the MA bit set to the Configurator node with the next lowest Unique ID. If no lower Unique ID exists, it sends a QUERY NODE SMS with the MA bit set to the Configurator node with the highest Unique ID.
- c) If 5 seconds elapses between receipt of QUERY NODE SMSs with the MA bit set, the Configurator node sends a QUERY NODE SMS with the MA bit set to the Master.
- d) If neither the QUERY NODE REPLY SMS to the QUERY NODE SMS with the MA bit set is sent to the Master nor the QUERY NODE SMS with the MA bit set is received within 5 seconds the Configurator node invokes the Configuration process.

13.3 SAT Region process

While it is optional for Configurators to create, extend or delete SAT regions, it is not optional for any Configurator capable of being a Master in a multi-Configurator web to reply to a REQUEST SAT REGION SMS. If the Configurator does not support the creation and management of SAT regions, it shall return a REQUEST SAT REGION REPLY SMS with a RETURN CODE field value of FAILURE - SAT REGIONS NOT SUPPORTED. The optional process that a Configurator uses to manage a SAT region is as follows consists of three parts, Create SAT region, Extend SAT region and Delete SAT region.

13.3.1 Create SAT region

To create a SAT region, a Configurator shall send a REQUEST SAT REGION SMS to the Master indicating the endpoints and SAT quotas of the requested region, as well as a SAT OPCODE field value of CREATE A SAT

REGION. The Master shall check its internal data structures for possible overlap of SAT regions, and violation of SAT region rules.

If an overlap exists, the Master shall return the region identifier of the overlapping region in the REGION IDENTIFIER field along with a RETURN CODE field value of FAILURE - REQUESTED CREATE OR EXTEND OVERLAPS EXISTING REGION in the REQUEST SAT REGION REPLY SMS. The requesting node may use this information to extend the existing region.

If there is no overlap, then the Master shall enter the new region's ENDPOINT #1 UNIQUE ID, ENDPOINT #2 UNIQUE ID, A QUOTA and B QUOTA values into its internal data structures. The Master shall return the region identifier of the region in the REGION IDENTIFIER field of the REQUEST SAT REGION REPLY SMS with a RETURN CODE field value of SUCCESSFUL COMPLETION.

If the creation was successful, the Master shall configure the ports in the region using the quotas supplied in the REQUEST SAT REGION SMS. The Master shall send a Master Alert with an ALERT CODE value of SAT MEMBERSHIP ESTABLISHED to each Configurator in the region. The endpoints are set to reflect SAT and SAT as necessary to form the region.

The endpoints of the region requested are defined by the ENDPOINT #1 UNIQUE ID and ENDPOINT #2 UNIQUE ID fields in conjunction with the ENDPOINT #1 PORT NUMBER and ENDPOINT #2 PORT NUMBER fields in the REQUEST SAT REGION SMS. If the region defined by the endpoints does not contain the requester, and the SAT OPCODE field in the REQUEST SAT REGION SMS is CREATE A SAT REGION or EXTEND A SAT REGION, then the Master returns a REQUEST SAT REGION REPLY SMS with a RETURN CODE value of WARNING THE REQUESTER THAT THE CREATED OR EXTENDED REGION DOES NOT INCLUDE THE REQUESTER. The REGION IDENTIFIER field is valid.

If the Master does not support the creation of SAT regions, then it returns a REQUEST SAT REGION REPLY SMS to the requester with a RETURN CODE of SAT REGIONS NOT SUPPORTED. The REGION IDENTIFIER field shall be ignored.

13.3.2 Extend SAT region

If a create region request resulted in a failure due to overlapped regions, then the requester may issue a REQUEST SAT REGION SMS with a SAT OPCODE field value of EXTEND SAT REGION. The ENDPOINT #1 UNIQUE ID, ENDPOINT #1 PORT NUMBER, ENDPOINT #2 UNIQUE ID, and the ENDPOINT #2 PORT NUMBER fields contain the new endpoint identifiers for the extended region. The quotas remain as set by the original creator, as does the value of the region identifier as assigned by the Master. The Master checks that the requested extend does not overlap another SAT region. The endpoints of the new region are the result of the union of the old set of nodes. Therefore the new region may have neither, one, or both endpoints moved from the original region. If there is no overlap caused by extending the region, then the Master returns a REQUEST SAT REGION REPLY SMS with a RETURN CODE of SUCCESSFUL OPERATION. The Master then configures the ports in the region using the quotas defined in the original create request. The Master also sends a Master Alert with an ALERT CODE value of SAT REGION MEMBERSHIP ESTABLISHED to all the Configurators that have been added to the region.

If overlap exists, the Master returns the region identifier of the overlapped region along with a RETURN CODE of FAILURE - REQUESTED CREATE OR EXTEND OVERLAPS EXISTING REGION. The REGION IDENTIFIER field of the REQUEST SAT REGION REPLY SMS contains the region identifier of the region that caused the conflict.

13.3.3 Delete SAT region

A Configurator may request the deletion of a region as follows in order:

- a) The Configurator creates a REQUEST SAT REGION SMS with an SAT OPCODE field of DELETE A SAT REGION. The node shall supply the region's region identifier in the REGION IDENTIFIER field. The ENDPOINT #2 UNIQUE ID, ENDPOINT #1 PORT NUMBER and ENDPOINT #2 PORT NUMBER fields are ignored.
- b) If the deletion is successful and if the region defined does not contain the requester, then the Master returns a REQUEST SAT REGION REPLY SMS with a status of SUCCESSFUL WITH WARNING: DELETED REGION DID NOT INCLUDE REQUESTER. The REGION IDENTIFIER field is valid.
- c) If the region defined does contain the requester, the Master returns a REQUEST SAT REGION REPLY SMS with a status of SUCCESSFUL OPERATION. The REGION IDENTIFIER field is valid.

- d) If the delete request was successful, the Master configures the ports to the system default state. The Master also sends a MASTER ALERT SMS with an ALERT CODE field value of SAT REGION MEMBERSHIP DELETED to all Configutors that reside in the region to be deleted. The Master then deletes the region from its internal data structures.
- e) If a new Master is elected, the SAT region information held by the old Master is lost. The new Master issues CONFIGURE PORT SMSs to all the ports in the system, using its system default values which effectively clears out all existing SAT regions. When the creators of SAT regions receive the CONFIGURE PORT SMS indicating a new Master, they may repeat the process of creating SAT regions, thus rebuilding the Master's SAT region database if the new Master supports SAT regions. If the endpoints, region identifier or port number fields in the REQUEST SAT REGION SMS are invalid, then the Master returns a REQUEST SAT REGION REPLY SMS with a status code of SMS REJECTED, INVALID ENDPOINT, REGION IDENTIFIER OR PORT NUMBER. Byte 4 contains the first byte number of the offending field in the request.

13.4 Configutor Table Maintenance process

The Master shall begin the Configutor Table Maintenance Process when either of the following events has occurred, and will perform the maintenance on the set of nodes indicated:

- a) A new Master takes over the Web, and process all Responders in the web.
- b) A new Sub-Web has been added to the web, and process the Responders in the new sub-web

This process is a low priority process that is run after the Master has built its Topology table of the web and after all ports have been placed in normal mode. The Master performs the following operations on each node in the system that has a Configutor table, and has come under the Master's control since the completion of its last Configutor Table Maintenance process:

- a) The Master sends a QUERY REGISTRATION SMS to the target node and waits for the reply.
- b) When the Master receives the associated Data Ready SMS, it reserves buffer space for as much of the Configutor table space as possible, and sets up a channel to receive the data. The Master reserves space for a clean-up list which will contain the Return Path IDs and port numbers for entries to be deleted.
- c) The Master sends a Data Reply SMS to the target providing it with the channel number and byte count for the transfer.
- d) When the Master receives the table data, it verifies that each Return Path ID points to the correct Unique ID node relative to the indicated port number. If it does not, then the Master adds the Return Path ID and port number to the clean-up list. The Responder sets a flag in each table entry indicating it has been sent to the Master in response to a QUERY REGISTRATION SMS. This flag is cleared if the table entry is subsequently changed via a quiesce, Total Reset or Absolute Reset frame.
- e) The Master sends additional Data Reply SMSs with offsets and byte counts necessary to complete the total table contents as indicated by the Data Ready SMS. On each portion of the table retrieved, step d is repeated.
- f) The Master sends DELETE RETURN PATH ID SMSs to the target node for each entry in the cleanup list.
- g) The Master waits for a RESPONSE SMS for each DELETE RETURN PATH ID SMS that it sent.

The Master sets an internal flag to indicate that the target node has been maintained since the Master took control of the node.

Annex A (Informative) Transport Layer Data Transfer Processes

A.1 Normal read transfer process

The TARGET READ, INITIATOR READ, and READ DATA COMPLETED services are for all or part of a normal read data transfer (see Annex B). The normal read data transfer consists of four parts as follows in order:

- a) The target s protocol layer issues a TARGET READ transport service that results in one or more DATA READY SMS(s) being sent (see B.4.4);
- b) The receipt of a DATA READY SMS triggers the initiator s protocol layer to invoke the INITIATOR READ transport service, that sends one or more DATA REPLY SMS(s) (see A.2);
- c) The receipt of a DATA REPLY SMS triggers the target s Send Read Data process (see A.2);
- d) Sending the last byte of data triggers the target s transport layer to invoke the Read Transfer Ended process (see A.3).

A.2 Send Read Data process

The Send Read Data process is a target only process that is triggered by the receipt of a DATA REPLY SMS. The DATA REPLY SMS is validated as described in 11.3.2 using the fields in the Outstanding Data Transfer Table entry pointed to by the TAG field of the DATA REPLY SMS. The transport layer then sends the read data as specified in the DATA REPLY SMS.

A.3 Read Transfer Ended process.

When the target completes a read data transfer on a given Channel it finds the entry in the Outstanding Data Transfers Table that matches the Channel. If the ARRAY OF CHANNEL field contains more than one element, then the transport layer clears that Channel element. If the ARRAY OF CHANNEL field contains only one entry, then the transport layer invokes the READ DATA COMPLETED transport service as follows and clears the entry in the Outstanding Data Transfers Table:

- a) The TAG field becomes the Tag parameter.
- b) The RETURN PATH ID field is used to reference the Configurator table to find the Initiator s Unique ID to return in the Unique ID parameter.
- c) The status of the data transfer determines the value of the Return Code parameter (FUNCTION SUCCESSFUL or FUNCTION FAILED).

A.4 Fast read transfer process

The TARGET FAST READ and FAST READ DATA COMPLETED services are intended for read data transfers where the initiator channel is already known, and no DATA READY or DATA REPLY SMSs need be exchanged. The service is initiated by the target and needs no initiator services, as the initiator has already set up for the data transfer. The fast read data transfer consists of two parts in order:

- a) The target s protocol layer issues a TARGET FAST READ transport service that results in read data being transmitted (see B.4.7).
- b) Sending the last byte of data triggers the target s transport layer to invoke the FAST READ DATA COMPLETED transport service (see B.4.8).

A.5 Write Transfer process

The TARGET WRITE, INITIATOR WRITE, and WRITE DATA COMPLETED services control normal write transfers. The normal write data transfer consists of three parts in order:

- a) The target's protocol layer issues a TARGET WRITE transport service that results in one or more DATA REQUEST SMS(s) being sent (see B.4.9);
- b) The receipt of a DATA REQUEST SMS triggers the initiator's protocol layer to invoke the INITIATOR WRITE transport service, that sends the write data (see B.4.1);
- c) Receiving the last byte of data triggers the target's transport layer to invoke the Write Transfer Ended transport process (see A.6).

A.6 Write Transfer Ended process.

When the target completes a write data transfer on a given Channel it finds the entry in the Outstanding Data Transfers Table that matches the Channel. If the ARRAY OF CHANNEL field contains more than one element, then the transport layer clears that Channel element. If the ARRAY OF CHANNEL field contains only one entry, then the transport layer invokes the WRITE DATA COMPLETED transport service as follows and clears the entry in the Outstanding Data Transfers Table:

- a) The TAG field becomes the Tag parameter;
- b) The RETURN PATH ID field is used to reference the Configurator table to find the Initiator's Unique ID to return in the Unique ID parameter;
- c) The status of the data transfer determines the value of the Return Code parameter (FUNCTION SUCCESSFUL or FUNCTION FAILED).

Annex B (Informative) Transport Layer Services

B.1 Introduction

The transport/protocol services are those services that define the logical boundary between the transport layer and the protocol layer. The services are not meant to define a hardware/software split, nor even a dividing line between software routines.

The service is either invoked by the Transport layer and execute in the Protocol layer, or invoked in the Protocol layer and execute in the Transport layer. Since this is the Transport layer, the appropriate clauses concentrates mainly on the transport layer s function in the service.

Table B.1 shows the services involved between the transport and protocol layers, along with who invokes them and where they execute.

Table B.1 - Transport services

Invoked by	Service name	Executes in	Initiator/target
Protocol	REGISTER ULP	Transport	Either
Transport	AER	Protocol	Either
Protocol	GET LOGICAL PATH	Transport	Either
Protocol	GET TAG	Transport	Either
Protocol	FREE TAG	Transport	Either
Protocol	GET UNIQUE ID	Transport	Either
Protocol	SEND SMS	Transport	Either
Transport	SMS RECEIVED	Protocol	Either
Protocol	TARGET READ	Transport	target
Transport	INITIATOR READ	Protocol	Initiator
Transport	READ DATA COMPLETED	Protocol	target
Protocol	TARGET FAST READ	Transport	target
Transport	FAST READ DATA COMPLETED	Protocol	target
Protocol	TARGET WRITE	Transport	target
Transport	INITIATOR WRITE	Protocol	Initiator
Transport	WRITE DATA COMPLETED	Protocol	target
Transport	QUIESCE	Protocol	target

The transport services are separated into six main groups.

B.1.1 SMS transfer

The GET LOGICAL PATH, GET TAG and SEND SMS transport service may be used to transmit an SMS frame. The SMS RECEIVED, GET UNIQUE ID, and FREE TAG transport services may be used when an SMS is received.

B.1.2 Normal read

The TARGET READ, INITIATOR READ, and READ DATA COMPLETED services are for all or part of a normal read data transfer.

B.1.3 Fast read

The TARGET FAST READ and FAST READ DATA COMPLETED services in intended for read data transfers where the initiator channel is already known, and no DATA READY or DATA REPLY SMSs need be exchanged. The service is initiated by the target and needs no initiator services, as the initiator has already set up for the data transfer.

B.1.4 Write

The TARGET WRITE, INITIATOR WRITE, and WRITE DATA COMPLETED services control write transfers.

B.1.5 Quiesce

The QUIESCE transport service is invoked when the target receives a QUIESCE SMS, that causes the target to invoke the Quiesce Protocol service to abort all related I/O processes. A RESPONSE SMS is then generated with a RETURN CODE field related to the Return Code parameter returned by the QUIESCE transport service.

B.1.6 Other

The REGISTER ULP service allows the ULP to register and specify its operation parameters to the transport layer. The AER service allows the transport layer to inform the ULP of asynchronous events.

B.2 Conventions

B.2.1 Naming conventions

The term Transport Services refers to those services between the SSA-TL2 transport layer and the protocol layer, and in following SSA-TL2 conventions the name is all letters capitalized.

The terms associated with the Service are called parameters and are first letter capitalized.

Any SMS fields referenced are indicated in all small capital letters.

B.2.2 Notations

The services are indicated with indented text as illustrated below:

Service-name (Input parm1, [Optional input parm2] || Output parm1, [Optional Output parm2])

The service name precedes the parenthesis. The parenthesis enclose the parameter lists. Any parameter enclosed in angle brackets ([...]) is considered optional. The double bar (||) separates the input parameters from the output parameters. The transport services have no optional parameters.

B.2.3 Unique ID vs. Return Path ID

The protocol layer uses a node's Unique ID to identify it. The concepts of Port or Path are unknown to the protocol layer. To generate an SMS with a RETURN PATH ID field, the protocol layer requests the transport layer to replace bytes 4 through 7 with the RETURN PATH ID associated with the Unique ID parameter using the Replace4-7 parameter. When an SMS is parsed and determined to contain a RETURN PATH ID field, the protocol layer determines the appropriate Unique ID by issuing the GET UNIQUE ID transport service.

The transport layer controls the Configurator table and Configuration table, and uses either Unique ID or RETURN PATH ID. The transport layer provides the conversion between RETURN PATH ID and Unique ID with the GET UNIQUE ID transport service. Since the transport layer controls the generation of RETURN PATH ID, it is indirectly responsible for the control of the use of ports and multiple paths.

B.3 Transport layer Data Structures

B.3.1 Configuration table

The Configuration table (see 9.7.3) resides in the Initiator, and is used to choose a port and path to a node when given its Unique ID, and to associate a Return Path ID with a Unique ID.

B.3.2 Configurator table

The Configurator table (see 9.7.4) resides in the Initiator, and is used to choose a port and path to a node when given its Unique ID, and to associate a Return Path ID with a Unique ID.

The Configurator table contains a deadlock flag for each node that controls the number of certain SMSs that are outstanding.

B.3.3 Outstanding Data Transfers Table

The target transport layer uses the Outstanding Data Transfers Table during data transfers and consists of an array of fields shown in Table B.2 for each outstanding data transfer. The initiator does not use this table.

Table B.2 - Outstanding Data Transfers Table entry

Field	Description
TAG	The Tag value supplied by the TARGET READ or TARGET WRITE transport services.
RETURN PATH ID	The Return Path ID chosen by the TARGET READ or TARGET WRITE transport services.
ARRAY OF CHANNEL	An array of Channel values used by the data transfer.
ARRAY OF TRANSFER ID	An array of Transfer ID values used by the data transfer.

When a TARGET READ or TARGET WRITE transport service is invoked, the transport layer constructs an entry in the Outstanding Data Transfers Table. The entry is cleared when the data transfer completes and the READ DATA COMPLETED or WRITE DATA COMPLETED transport service is invoked.

The TAG field is loaded directly from the Tag parameter of the associated transport service. The TAG field is used to construct SMSs, identify a transfer with an I/O process, and to return the Tag parameter of the READ DATA COMPLETED and WRITE DATA COMPLETED transport services.

The RETURN PATH ID field is loaded from the Return Path ID that was chosen by the transport layer to send the DATA READY or DATA REQUEST SMS. The RETURN PATH ID field is used to construct the Unique ID parameter of the READ DATA COMPLETED and WRITE DATA COMPLETED transfer services. The RETURN PATH ID field is also used to validate DATA REPLY SMSs.

The ARRAY OF CHANNEL field is an array of Channel values being used by the data transfer. An entry is made each time the node is set up to accept a Channel, and is cleared when the Channel is no longer active for that data transfer. When a data transfer completes, the table is searched by Channel to find the Tag and Unique ID to be reported in the READ DATA COMPLETE or WRITE DATA COMPLETE transport service.

The ARRAY OF TRANSFER ID field is an array of Transfer ID values being used by a read data transfer. An entry is made whenever a DATA READY SMS is sent, and is cleared when the last DATA REPLY SMS associated with the DATA READY SMS is received and validated.

B.3.4 Object definitions

The parameters for the Transport/Protocol services are as follows:

- a) The Buffer Address parameter is a memory location where data is to be stored to or retrieved from.
- b) The Byte Offset parameter is a 4 byte value indicating the offset of this data transfer based on the Buffer Address parameter.
- c) The Byte Count parameter is a 4 byte value indicating the number of bytes for a data transfer. In the case of a split read or write, the byte count field only indicates the amount of data to be transferred during this portion of the split.
- d) The Control parameter indicates the CONTROL field for or from the SMS frame. It consists of the FRAME TYPE field and optionally the RESET TYPE field (if FRAME TYPE value of CONTROL FRAME). The FSN is not part of the CONTROL parameter, and is zero filled.
- e) The Length parameter is a 1 byte count of the number of bytes in the DATA field of the frame to be sent or received.
- f) The Channel parameter is a 2 byte value indicating the initiator channel for the target to use as the channel component in the ADDRESS field of the read Data frame(s).
- g) The Return Path ID parameter is the 4 byte value used by some SMSs to identify a return path.
- h) The In Order parameter is a 1 bit flag that indicates that the data is written into the buffer space as received (In Order set to 0b), or written into the buffer space in order with the Buffer Address location containing the byte of data with an offset of zero (In order set to 1b). The In Order parameter allows the transport layer to use split data transfers even if the Protocol layer does not support splits.

- i) The Replace4-7 parameter is a flag to indicate whether the transport layer replaces bytes 4 through 7 of the SMS with the Return Path ID associated with the Unique ID parameter.
- j) The Return Code parameter is an indicator of the success or failure of the service to execute properly. The valid value of the Return Code parameter are shown in Table B.3.
- k) The Start Count parameter is a 4 byte value indicating the byte count to be used for the first DATA READY, DATA REPLY, or DATA REQUEST SMS.
- l) The Tag parameter is a 2 byte value corresponding to the Tag value of the I/O process.
- m) The Channel parameter is a 2 byte value indicating the target channel for the initiator to use as the channel component in the ADDRESS field of the write Data frame(s).
- n) The Threshold Count parameter is a 4 byte value indicating the amount of available buffer space that triggers another DATA REPLY or DATA REQUEST SMS, or the amount of data available that triggers the next DATA READY SMS.
- o) The Unique ID parameter is the 8 byte value reported by a node in the UNIQUE ID field of the QUERY NODE REPLY SMS and is a globally unique identifier of a node. Each Configurator in the Web uses the same Unique ID to identify a node.

Table B.3 - Return Code parameter values

Value	Name	Meaning
00h	FUNCTION SUCCESSFUL	The requested function completed successfully
01h-FDh	reserved	
FEh	FUNCTION FAILED	The requested function failed
FFh	INVALID PARAMETER	One or more parameters or fields were invalid

B.4 The Transport Services

B.4.1 GET LOGICAL PATH

GET LOGICAL PATH (Destination ID, Logical Path Number || Return Path ID, Return Code)

The GET LOGICAL PATH initiator transport service request is used during the construction of a protocol layer SMS to assign a Return Path ID for a given Logical Path Number. The input parameters are the Destination ID and Logical Path Number. The Output parameters are Return Path ID and Return Code.

The Destination ID parameter is the identifier for the Destination node.

The Logical Path Number is a number for a logical path. The same Return Path ID will always be returned for a given Logical Path Number.

The Return Path ID parameter is the Return Path ID to be used for that Logical Path. The same Return Path ID may be used for multiple Logical Path Numbers, but a given Logical Path Number will always return the same Return Path ID.

The Return Code parameter indicates the status of the GET LOGICAL PATH service request. A value of INVALID PARAMETER indicates that the Destination ID is not known.

B.4.2 GET TAG

GET TAG (Destination ID || Tag, Return Code)

The GET TAG initiator transport service request is used during the construction of a protocol layer SMS to assign a Tag value. The input parameter is the Destination ID. The Output parameters are Tag and Return Code. This service request does not need to exist and Tag assignment may be controlled within the protocol layer if only one protocol layer is running in the initiator at a time. The GET TAG service request needs to be implemented only in initiators that have multiple protocols running at any given time.

The Destination ID parameter is the identifier for the Destination node.

The Tag parameter is the value to be used for the TAG field, and is unique for all outstanding Tag values to that node. The Tag value is freed with a FREE TAG transport layer service request, and is only valid if the Return Code parameter indicates FUNCTION SUCCESSFUL.

The Return Code parameter indicates the status of the GET TAG service request. A value of FUNCTION FAILED indicates that no unique Tag value is available at this time.

B.4.3 FREE TAG

FREE TAG (Destination ID, Tag || Return Code)

The FREE TAG initiator transport service request is used during the construction of a protocol layer SMS to free a Tag value. The input parameters are the Destination ID and Tag. The Output parameter is Return Code. This service request does not need to exist and Tag assignment may be controlled within the protocol layer if only one protocol layer is running in the initiator at a time. The FREE TAG service request needs to be implemented only in initiators that have multiple protocols running at any given time.

The Destination ID parameter is the identifier for the Destination node.

The Tag parameter is the value to be freed. The Tag value may be considered freed if the Return Code parameter indicates FUNCTION SUCCESSFUL.

The Return Code parameter indicates the status of the FREE TAG service request. A value of INVALID PARAMETER indicates that the Destination ID is not known, or the Tag value was not assigned.

B.4.4 GET UNIQUE ID

GET UNIQUE ID (Return Path ID || Unique ID, Return Code)

The GET UNIQUE ID transport service has one input parameter (Return Path ID) and two output parameters (Unique ID and Return Code). It is initiated in the protocol layer and executes in the transport layer.

The Unique ID is obtained by the Transport layer by scanning the Return Path ID array elements of the Configurator table or Configuration table and returning the associated Unique ID.

A Return Code parameter value of FUNCTION SUCCESSFUL validates the Unique ID parameter. A Return Code parameter value of FUNCTION FAILED, indicates the Return Path ID is not in the Configurator table.

B.4.5 SEND SMS

SEND SMS (Control, Unique ID, Buffer Address, Length, Replace4-7 || Return Code)

The SEND SMS transport service input parameters are Control, Unique ID, Buffer Address, Length, and Replace4-7. The return parameter is Return Code. It is initiated by the protocol layer and runs in the transport layer.

The SEND SMS transport service requests the transport layer to create and send a frame specified by the input parameters as described in the following:

- a) The CONTROL field is the control parameter with the addition of the FSN that is added by the transport layer hardware.
- b) The ADDRESS field of the frame is constructed by the Transport layer by using the path within the Configuration table (initiators) or Configurator table (targets) pointed to by the Unique ID parameter, combined with a Channel of 00h. The Unique ID parameter is the destination's Unique ID. The choice of the path and port is controlled by the transport layer and is vendor specific.
- c) The DATA field is the number of bytes specified by the Length parameter beginning at the address pointed to by the Buffer Address pointer.
- d) If the Replace4-7 parameter is set, then the transport layer replaces bytes 4 through 7 of the SMS with the RETURN PATH ID field from the Configuration table. The Replace4-7 parameter is only valid for initiator transport layers.

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the proper sending of the frame (ACK received), and not its successful receipt and parsing. A recovered error, such as a successful retry during Link ERP is considered successful and results in a FUNCTION SUCCESSFUL Return Code value. A Return Code

parameter value of FUNCTION FAILED is used when the node was unable to send the SMS, such as not having that Unique ID in the Initiator or Configuration tables, or the appropriate port is non operational.

B.4.6 SMS RECEIVED

SMS RECEIVED (Control, Buffer Address, Length)

The SMS RECEIVED transport service input parameter are Control, Buffer Address, and Length. There are no output parameters. It is initiated by the Transport layer when an SMS is received over the link and passes the Transport layer SMS Validation process (see 11.1). The SMS RECEIVED process executes in the Protocol layer.

The SMS RECEIVED transport service supplies the Protocol layer with the information it needs to process the SMS as follows:

- a) The CONTROL field is the control parameter.
- b) The DATA field is the number of bytes specified by the Length parameter beginning at the address pointed to by the Buffer Address pointer.
- c) The length of the DATA field is the Length parameter.

Any errors in receiving an SMS are handled by the transport layer and are not surfaced to the protocol layer (i.e., they do not cause the SMS RECEIVED service to be invoked).

B.4.7 TARGET READ

TARGET READ (Unique ID, Tag, Buffer Address, Byte Offset, Byte Count, Start Count, Threshold Count | | Return Code)

The TARGET READ transport service input parameters are Unique ID, Tag, Buffer Address, Byte Offset, Byte Count, Start Count, Threshold Count. The return parameters are Return Code. It is invoked by the Protocol layer and executes in the Transport layer.

The TARGET READ transport service instructs the transport layer to initiate a data transfer by generating a DATA READY SMS as follows:

- a) The CONTROL field of the frame is 00h.
- b) The ADDRESS field of the frame is constructed by the Transport layer by using the path within the Configurator table pointed to by the Unique ID parameter, combined with a Channel of 00h. The choice of the path and port is controlled by the transport layer and is vendor specific.
- c) The SMS CODE field is 20h.
- d) The TAG, BYTE OFFSET, and START COUNT fields are the Tag, Byte Offset, and Byte Count parameters respectively.
- e) The TRANSFER ID field is controlled by the Transport layer and varies with the number of outstanding data transfers.

If the Start Count parameter is less than the Byte Count parameter, then additional DATA READY SMSs are generated according to the rules specified in 11.3.1 as follows in order:

- a) Wait until the buffer has the amount of data specified in the Threshold Count parameter.
- b) The Byte Count parameter is set to the greater of zero or the Byte Count parameter decremented by the Start Count parameter.
- c) The Start Count parameter is set to the lesser of the Byte Count or Threshold Count parameters.
- d) A DATA READY SMS is generated as described previously in this clause.

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the DATA READY SMSs were sent successfully. A Return Code parameter value of FUNCTION FAILED indicates the SMS transfer was incomplete. Possible reasons for being incomplete are Link ERP exit errors, or parameter errors.

An entry is made in the Outstanding Data Transfers Table consisting of the Tag parameter, The Return Path ID chosen by the transport layer (based on the Unique ID), all the Channel values chosen by the transport layer, and all the Transfer ID values chosen by the transport layer.

B.4.8 INITIATOR READ

INITIATOR READ (Tag, Byte Offset, Byte Count || Return Path ID, Buffer Address, Start Count, Threshold Count, In Order, Return Code)

The INITIATOR READ service input parameters are Tag, Byte Offset, and Byte Count. The output parameters are Buffer Address, Start Count, Threshold Count, In Order, and Return Code. The service is invoked by the transport layer when a DATA READY SMS is received, and runs in the protocol layer.

The INITIATOR READ service requests the initiator protocol layer to allocate buffer space for the read data, and prepare to receive the data. The initiator allocates buffer space based on the Tag, Byte Offset, and Byte Count. The initiator protocol layer then returns the buffer segment length (Start Length), the amount of buffer that shall be empty before the service issues another DATA REPLY SMS (Threshold Length), a flag to tell the service to store the data in order in the buffer (In Order), and a Return Code.

A Return Code value of FUNCTION SUCCESSFUL instructs the service to set up the hardware to transfer the data using the Return Path ID to determine the path, and the Channel from the DATA REQUEST SMS. A Return Code value of FUNCTION FAILED instructs the service to ignore the DATA REQUEST SMS and transmit no data. After returning a FUNCTION FAILED Return Code, it is the responsibility of the protocol layer to terminate the I/O process.

The INITIATOR READ transport service instructs the transport layer to set up hardware for a data transfer and generate a DATA REPLY SMS as follows:

- a) The CONTROL field of the frame is 00h.
- b) The ADDRESS field of the frame is constructed by the Transport layer by using the path and port within the Configurator table pointed to by the Return Path ID parameter, combined with a Channel of 00h. The choice of the path and port is controlled by the transport layer and is vendor specific.
- c) The SMS CODE field is 21h.
- d) The RETURN PATH ID field is the Return Path ID parameter.
- e) The BYTE COUNT field is the Start Count parameter.
- f) The CHANNEL field is controlled by the transport layer based on what buffer pointer set was used to receive the read data.
- g) The TRANSFER ID field from the TRANSFER ID field on the DATA READY SMS.

If the Start Count parameter is less than the Byte Count parameter, then additional DATA REPLY SMSs are generated according to the rules specified in 11.3.1 as follows in order:

- a) Wait until buffer has the amount of space indicated by the Threshold Count parameter.
- b) The Byte Count parameter is set to the greater of zero or the Byte Count parameter decremented by the Start Count parameter.
- c) The Start Count parameter is set to the lesser of the Byte Count or Threshold count parameters.
- d) A DATA REPLY SMS is generated as described previously in this clause.

If the In Order parameter is set, the transport layer arranges for the data to appear in the buffer in the order it was originally written to the target. This means that the data is stored beginning at the Start Address plus the Byte Offset parameter values.

B.4.9 READ DATA COMPLETED

READ DATA COMPLETED (|| Tag, Unique ID, LUN, Return Code)

The READ DATA COMPLETED service has no input parameters and has four output parameters (Tag, Unique ID, LUN, and Return Code).

The READ DATA COMPLETED service is initiated by the Read Transfer Ended transport process. The Tag, Unique ID, and LUN parameters identify the read data transfer that completed.

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the proper sending of the read Data frames (ACK received), and not its successful receipt and parsing. A recovered error, such as a successful retry during Link ERP is considered successful and results in a FUNCTION SUCCESSFUL Return Code value. A Return Code parameter value of FUNCTION FAILED is used when the node was unable to send the SMS,

such as not having that Unique ID in the Initiator or Configuration tables, or the appropriate port is non Operational.

B.4.10 TARGET FAST READ

TARGET FAST READ (Unique ID, Byte Count, Channel || Return Code)

The TARGET FAST READ transport service input parameters are Unique ID, Byte Count, and Channel. The return parameter is Return Code.

The service is invoked by the protocol layer and executes in the transport layer. The TARGET FAST READ transport service then constructs a set of Data frames (128 bytes maximum each) to transfer the number of bytes in the Byte Count parameter as follows:

- a) The CONTROL field of the frame is 00h;
- b) The ADDRESS field of the frame is constructed by the Transport layer by using the path within the Configurator table pointed to by the Unique ID parameter, combined with a Channel of the Channel parameter.

B.4.11 FAST READ DATA COMPLETED

FAST READ DATA COMPLETED (|| Channel, Unique ID, Return Code)

The READ DATA COMPLETED service has no input parameters and has two output parameters (Channel and Return Code).

The READ DATA COMPLETED service is initiated by the transport layer when the data transfer is concluded for the operation identified by the Channel and Unique ID parameters.

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the proper sending of the read Data frames (ACK received), and not its successful receipt and parsing. A recovered error, such as a successful retry during Link ERP is considered successful and results in a FUNCTION SUCCESSFUL Return Code value. A Return Code parameter value of FUNCTION FAILED is used when the node was unable to send the SMS, such as not having that Unique ID in the Initiator or Configuration tables, or the appropriate port is non Operational.

B.4.12 TARGET WRITE

TARGET WRITE (Unique ID, Tag, Buffer Address, Byte Offset, Byte Count || Return Code)

The TARGET WRITE transport service input parameters are Unique ID, Tag, Buffer Address, Byte Offset, and Byte Count. The return parameters are Return Code. The service is invoked by the protocol layer and executes in the transport layer.

The TARGET WRITE transport service sets up the node to accept the number of bytes of data specified by the Byte Count parameter to be stored in the buffer starting at the address specified in the Buffer Address parameter. It assigns a Channel value to be used in the DATA READY SMS and to route the data.

The TARGET WRITE transport service constructs a DATA READY SMS as follows:

- a) The CONTROL field of the frame is 00h.
- b) The ADDRESS field of the frame is constructed by the transport layer by using the path component from the Configurator table pointed to by the Unique ID parameter, and the channel component of 00h. The choice of the path and port is controlled by the transport layer and is vendor specific.
- c) The SMS CODE field is 22h.
- d) The TAG field is the Tag parameter.
- e) The BYTE OFFSET field is the Byte Offset parameter.
- f) The BYTE COUNT field is the Byte Count parameter.
- g) The CHANNEL field is controlled by the transport layer (depends on the number of outstanding data transfers and the number of channels supported by the target).

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the DATA REQUEST SMS was sent successfully. A Return Code parameter value of FUNCTION FAILED indicates the SMS transfer was incomplete. Possible reasons for being incomplete are Link ERP exit errors, or parameter errors.

An entry is made in the Outstanding Data Transfers Table consisting of the Tag parameter, The Return Path ID chosen by the transport layer (based on the Unique ID), all the Channel values chosen by the transport layer, and all the Transfer ID values chosen by the transport layer.

B.4.13 INITIATOR WRITE

INITIATOR WRITE (Tag, Byte Count, Byte Offset || Unique ID, Buffer Address, Return Code)

The INITIATOR WRITE transport service input parameters are Tag, Byte Count, and Byte Offset. The return parameters are Buffer Address and Return Code. The service is invoked by the transport layer and executes in the protocol layer.

The INITIATOR WRITE transport service is invoked by the transport layer with the receipt of a DATA REQUEST SMS, and requests the initiator protocol layer to instruct the physical layer as to the buffer location of the data. If a buffer underrun condition occurs, the transport layer does not send any frames until sufficient data is available.

Upon receipt of a DATA REQUEST SMS, the transport layer loads the TAG, BYTE COUNT and BYTE OFFSET fields of the DATA REQUEST SMS into the Tag Byte Count, and Byte Offset parameters and invokes the INITIATOR WRITE transport service.

A Return Code parameter value of FUNCTION SUCCESSFUL instructs the transport service to set up the node to send the number of bytes of data specified by the Byte Count parameter from the buffer starting at the address specified in the Buffer Address parameter. The Data frames shall be constructed as follows:

- a) The CONTROL field of the frame is 00h.
- b) The ADDRESS field of the frame is constructed by the transport layer by using the path component from the Configurator table pointed to by the Unique ID parameter, and the channel component of the value specified in the CHANNEL field of the DATA REQUEST SMS. The choice of the path and port is controlled by the transport layer and is vendor specific.
- c) The DATA field is constructed from the data being sent.

A RETURN CODE value of FUNCTION FAILED instructs the service to ignore the DATA REQUEST SMS and transmit no data. After returning a FUNCTION FAILED Return Code, it is the responsibility of the protocol layer to terminate the I/O process.

B.4.14 WRITE DATA COMPLETED

WRITE DATA COMPLETED (|| Tag, Unique ID, LUN, Return Code)

The WRITE DATA COMPLETED service has no input parameters and has four output parameters (Tag, Unique ID, LUN, and Return Code).

The WRITE DATA COMPLETED service is initiated by the Write Transfer Ended transport process. The Tag, Unique ID, and LUN parameters identify the read data transfer that completed.

A Return Code parameter value of FUNCTION SUCCESSFUL indicates the proper receipt of the write Data frames. A recovered error, such as a successful retry during Link ERP is considered successful and results in a FUNCTION SUCCESSFUL Return Code value. A Return Code parameter value of FUNCTION FAILED is used when an error occurred during data transfer such as too much data received, or the Link ERP failed.

B.4.15 QUIESCE

QUIESCE (Source ID || Return Code)

The QUIESCE transport service input parameter is Source ID. The return parameter is Return Code. The service is invoked by the transport layer and executes in the protocol layer.

The QUIESCE transport service is invoked by the transport layer with the receipt of a QUIESCE SMS, and instructs the Initiator to terminate all I/O processes issued by the return path identified by the Source ID parameter. The transport layer then generates a RESPONSE MESSAGE as follows:

- a) The CONTROL field of the frame is 00h.
- b) The ADDRESS field of the frame is constructed by the transport layer by using the path component from the RETURN PATH field from the QUIESCE SMS, and the channel component of 00h. The port is the same port that received the QUIESCE SMS.
- c) The SMS CODE field is 03h.
- d) The RETURN CODE field is the Return Code parameter.

B.4.16 REGISTER ULP

REGISTER ULP (Vendor Code, AER Mask || ULP Code, Return Code)

The REGISTER ULP service has two input parameters (Vendor Code and AER Mask) and two output parameters (ULP Code and Return Code).

The Register ULP service is invoked by the protocol layer when it is initialized. The protocol layer passes its Vendor Code that contains a single byte ULP code or a 4 byte protocol identifier (see 12.2.10), and an AER Mask that identifies the asynchronous events for which it wishes to be informed. The AER mask is a two byte construct consisting of a bit significant map to enable each asynchronous event (see Table B.4).

Table B.4 - AER Mask

Byte	Bit 7	6	5	4	3	2	1	Bit 0
MSB	AEE15	AEE14	AEE13	AEE12	AEE11	AEE10	AEE9	AEE8
LSB	AEE7	AEE6	AEE5	AEE4	AEE3	AEE2	AEE1	AEE0

The transport layer will return a single byte ULP Code that the protocol uses in its SMSs, and a Return Code. A Return Code value of FUNCTION SUCCESSFUL indicates proper registration of the protocol and validates ULP Code. A Return Code value of FUNCTION FAILED, indicates that the transport layer was unable to register the ULP because either too many vendor unique protocols are already assigned, or some other internal constraint, and the ULP Code is invalid.

When the transport layer completes a Register ULP service, the transport layer invokes the AER service with a meaning code of NEW DEVICE ADDED for each device in its Configuration table that supports the registered ULP. The transport layer then invokes the Asynchronous Alert process with an ALERT CODE field value of SUPPORTED ULP LIST CHANGED.

B.4.17 AER

AER (AER Code, Device ID, Alert Code, Frame Data ||)

The AER service has four input parameters (AER Code, Device ID, Alert Code, and Frame Data) and no output parameters. The AER service is a high priority service and should take precedence over other services.

The AER service is invoked by the transport layer when a Total Reset is processed, certain Asynchronous Alerts are invoked, or certain Master Alerts are received; and the appropriate bit in the AER Mask (from the REGISTER ULP service) is set. The valid AER Code values are shown in Table B.5.

Table B.5 - AER Code values

AER Code value	AER Mask bit	Meaning
0	AEE0	LOCAL RESET PROCESSED
1	AEE1	TOTAL RESET PROCESSED
2	AEE2	reserved
3	AEE3	reserved
4	AEE4	NEW DEVICE ADDED
5	AEE5	DEVICE ACCESS LOST
6	AEE6	DEVICE PATH CHANGED
7	AEE7	reserved
8	AEE8	INVALID SMS
9	AEE9	ENVIRONMENTAL ERROR
10	AEE10	WARNING
11	AEE11	reserved
12	AEE12	reserved
13	AEE13	reserved
14	AEE14	reserved
15	AEE15	reserved

An AER Code parameter value of LOCAL RESET PROCESSED indicates that the transport layer has been reset by an internal Local Reset. A Local Reset is equivalent to a Total Reset, but is invoked from within the node rather than as a frame over the SSA interface. The Device ID, Alert Code and Frame Data parameters are not valid.

An AER Code parameter value of TOTAL RESET PROCESSED indicates that the transport layer has been reset by an internal Local Reset. The Device ID, Alert Code and Frame Data parameters are not valid. The receipt of an Absolute Reset will be treated like a power on situation, and both transport and protocol layers will be automatically reset.

An AER Code parameter value of NEW DEVICE ADDED indicates that the transport layer has detected that a new device the supports the ULP has been added to the Web. The Device ID parameter is the new Device s 8 byte Unique ID. The Alert Code and Frame Data parameters are not valid.

An AER Code parameter value of DEVICE ACCESS LOST indicates that the transport layer has lost all paths to a device the supports the ULP. The Device ID parameter is the lost Device s 8 byte Unique ID. The Alert Code and Frame Data parameters are not valid.

An AER Code parameter value of DEVICE PATH CHANGED indicates that the transport layer has detected that the path has changed to a device the supports the ULP. The Device ID parameter is the Device s 8 byte Unique ID. The Alert Code and Frame Data parameters are not valid.

An AER Code value of INVALID SMS indicates that the transport layer has received a MASTER ALERT SMS with an ALERT CODE value of UNKNOWN SMS, SMS TOO SHORT, or INVALID FIELD. The Device ID parameter is the Device s 8 byte Unique ID who detected the error. The Alert Code parameter is the ALERT CODE value from the MASTER ALERT SMS. The Frame Data parameter contains the FRAME DATA field from the MASTER ALERT SMS.

An AER Code value of ENVIRONMENTAL ERROR indicates that the transport layer has invoked an Asynchronous Alert or received a MASTER ALERT SMS with an ALERT CODE value of POWER FAULT or CRITICAL FAN FAILURE. The Device ID parameter is 8 byte Unique ID of the device that detected the error. The Alert Code parameter is the ALERT CODE value from the Asynchronous Alert or MASTER ALERT SMS. The Frame Data parameter is not valid.

An AER Code value of WARNING indicates that the transport layer has invoked an Asynchronous Alert or received a MASTER ALERT SMS with an ALERT CODE value of REDUNDANT FAN FAILURE, FAILURE PREDICTION THRESHOLD EXCEEDED, or ALARM THRESHOLD EXCEEDED. The Device ID parameter is 8 byte Unique ID of the device that detected the error. The Alert Code parameter is the ALERT CODE value from the Asynchronous Alert or MASTER ALERT SMS. The Frame Data parameter is not valid.

Annex C (Informative) Examples of half and full duplex

The following examples illustrate the link protocol. For compactness, the diagrams show a 2-byte ADDRESS field and a 4-byte DATA field. Each port has a pair of A/B transmit frame buffers and a pair of A/B receive frame buffers. It is also assumed that the source/destination process is fast enough to fill/empty one buffer while the link empties/fills the other. Thus frames may be transmitted back-to-back, without any intervening delay.

C.1 Half Duplex

A half-duplex transfer with port 1 acting as a source and port 2 acting as a destination is shown in Figure C.1. Port 2 sends an RR character pair as soon as it detects the start of a frame since its receiver has buffer space for another frame. This clears the WAITING FOR RR FLAG in port 1 and allows another frame to be sent immediately after the trailing FLAG of the current frame. In addition, port 2 issues an ACK character pair immediately when it detects the trailing FLAG of each frame.

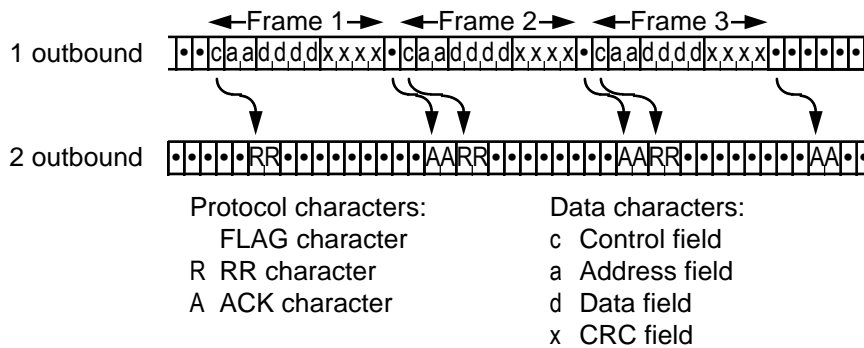


Figure C.1 - Half-duplex transfer

C.2 Full Duplex

A full-duplex transfer with port 1 acting as a source of 3 frames and port 2 acting as a source of 2 frames is shown in Figure C.2. This illustrates how responses are interleaved within a frame.

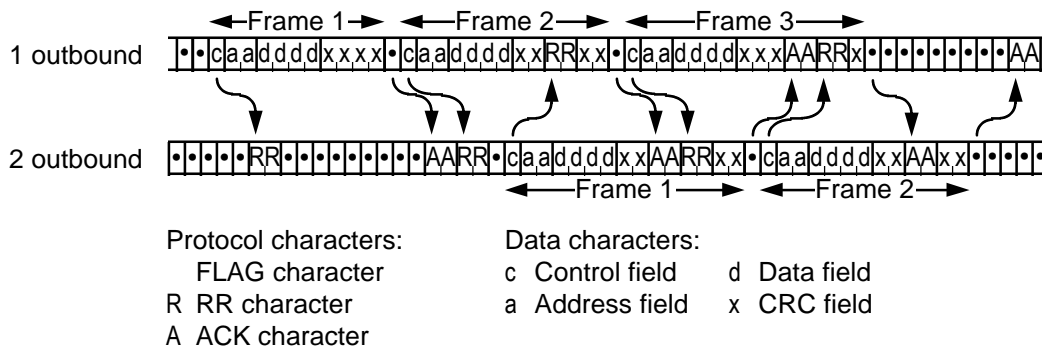


Figure C.2 - Full-duplex transfer