

Working Draft

**T10
Project 1155D**

**Revision 3
November 10, 1997**

Information technology — Serial Bus Protocol 2 (SBP-2)

This is a draft proposed American National Standard under development by T10, a Technical Committee of the National Committee for Information Technology Standardization (NCITS). As such, this is not a completed standard and has not been approved. The Technical Committee may modify this document as a result of comments received during public review and its approval as a standard.

Permission is granted to members of 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 commercial or for-profit replication or republication is prohibited.

T10 Technical Editor:

Peter Johansson
Congruent Software, Inc.
3998 Whittle Avenue
Oakland, CA 94602
USA

(510) 531-5472
(510) 531-2942 FAX

pjohansson@aol.com

Reference numbers
ISO/IEC xxxxx:199x
ANSI NCITS.xxx-199x

Printed November 10, 1997

Points of contact

T10 Chair:

John B. Lohmeyer
Symbios Logic, Inc.
4420 Arrows West Drive
Colorado Springs, CO 80907
USA

(719) 533-7560
(719) 533-7036 FAX
john.lohmeyer@symbios.com

T10 Vice-chair:

Lawrence J. Lamers
Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, CA 95035

(408) 957-7817
(408) 957-7193 FAX
ljlammers@aol.com

NCITS Secretariat:

NCITS Secretariat
1250 I Street NW, Suite 200
Washington, DC 2000
USA

(202) 737-8888
(202) 638-4922 FAX

T10 Bulletin board:

(719) 533-7950

T10 FTP:

[ftp.symbios.com/pub/standards/io/t10](ftp:symbios.com/pub/standards/io/t10)

T10 Home page:

<http://www.symbios.com/t10>

T10 Reflector:

scsi@symbios.com
majordomo@symbios.com (to subscribe)

IEEE 1394 Reflector:

p1394@sun.com
bob.snively@sun.com (to subscribe)

Document distribution:

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

(800) 854-7179
(303) 792-2181
(303) 792-2192 FAX

American National Standard
for Information Systems –

Serial Bus Protocol 2 (SBP-2)

Secretariat

Information Technology Industry Council

Not yet approved

American National Standards Institute, Inc.

Abstract

This standard specifies a protocol for the transport of commands, data and status between devices connected by Serial Bus, a memory-mapped split-transaction bus defined by IEEE Std 1394-1995. In order to take advantage of unique capabilities of Serial Bus for the transport of isochronous data, this standard provides methods to manage isochronous connections and to control the flow of isochronous data between devices.

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 effort be made towards their resolution.

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

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

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

Published by

**American National Standards Institute
1430 Broadway, New York, NY 10018**

Copyright © 1996, 1997 by American National Standards Institute
All rights reserved.

Printed in the United States of America

Contents

	Page
Foreword.....	vii
Revision history.....	ix
1 Scope and purpose.....	1
1.1 Scope	1
1.2 Purpose.....	1
2 Normative references	3
2.1 Approved references.....	3
2.2 References under development.....	3
3 Definitions and notation	5
3.1 Definitions	5
3.1.1 Conformance.....	5
3.1.2 Glossary	5
3.1.3 Abbreviations.....	8
3.2 Notation.....	8
3.2.1 Numeric values.....	8
3.2.2 Bit, byte and quadlet ordering.....	9
3.2.3 Register specifications	10
3.2.4 State machines	11
4 Model (informative).....	13
4.1 Unit architecture.....	13
4.2 Logical units	13
4.3 Requests and responses	13
4.4 Data buffers.....	14
4.5 Target agents	16
4.6 Ordered and unordered execution	17
4.7 Streams.....	17
4.7.1 Stream task set.....	19
4.7.2 Stream controller	20
4.7.3 Error reporting	20
5 Data structures	23
5.1 Operation request blocks (ORB's)	23
5.1.1 Dummy ORB	25
5.1.2 Command block ORB's.....	25
5.1.3 Stream control ORB	28
5.1.4 Management ORB's.....	32
5.2 Page tables	42
5.2.1 Unrestricted page tables.....	42
5.2.2 Normalized page tables.....	43
5.3 Status block.....	44
6 Control and status registers.....	49
6.1 Core registers	49
6.2 Serial Bus-dependent registers.....	49
6.3 MANAGEMENT_AGENT register.....	50
6.4 Command block and stream control agent registers	51
6.4.1 AGENT_STATE register.....	51
6.4.2 AGENT_RESET register	52

6.4.3 ORB_POINTER register.....	53
6.4.4 DOORBELL register.....	54
6.4.5 UNSOLICITED_STATUS_ENABLE register	54
7 Configuration ROM.....	55
7.1 Power reset initialization.....	55
7.2 Bus information block.....	55
7.3 Root directory.....	57
7.3.1 Module_Vendor_ID entry.....	57
7.3.2 Node_Capabilities entry.....	57
7.3.3 Unit_Directory entry.....	57
7.4 Unit directory	58
7.4.1 Unit_Spec_ID entry.....	58
7.4.2 Unit_SW_Version entry	58
7.4.3 Command_Set_Spec_ID entry.....	59
7.4.4 Command_Set entry	59
7.4.5 Command_Set_Revision entry.....	59
7.4.6 Firmware_Revision entry.....	60
7.4.7 Management_Agent entry	60
7.4.8 Logical_Unit_Characteristics entry.....	60
7.4.9 Logical_Unit_Directory entry.....	61
7.4.10 Logical_Unit_Number entry.....	61
7.4.11 Unit_Unique_ID entry.....	62
7.5 Logical unit directory	62
7.5.1 Command_Set_Spec_ID entry.....	62
7.5.2 Command_Set entry	63
7.5.3 Command_Set_Revision entry.....	63
7.5.4 Logical_Unit_Characteristics entry.....	63
7.5.5 Logical_Unit_Number entry.....	63
7.6 Unit unique ID leaf.....	63
8 Access	65
8.1 Access protocols.....	65
8.2 Login requests.....	65
8.2.1 Login.....	65
8.2.2 Create stream.....	66
8.3 Reconnection	67
8.4 Logout	68
9 Command execution.....	69
9.1 Requests and request lists.....	69
9.1.1 Fetch agent initialization (informative)	69
9.1.2 Dynamic appends to request lists (informative)	70
9.1.3 Fetch agent use by the BIOS (informative)	71
9.1.4 Fetch agent state machine	71
9.2 Data transfer	74
9.3 Completion status	74
9.4 Unsolicited status	75
10 Task management	77
10.1 Task sets	77
10.2 Basic task management model.....	77
10.3 Error conditions	78
10.4 Task management requests.....	78
10.4.1 Abort task	78
10.4.2 Abort task set	79

10.4.3 Clear task set	80
10.4.4 Logical unit reset	81
10.4.5 Target reset	81
10.4.6 Terminate task	82
10.5 Task management event matrix.....	82
11 Isochronous data interchange format	85
11.1 Cycle marks.....	85
11.2 Isochronous data packets	85
11.3 Null packets.....	86
11.4 Cycle mark index.....	87
11.5 Common isochronous packets (CIP)	87
12 Isochronous operations	91
12.1 Stream command block requests	91
12.2 Stream control.....	92
12.2.1 Channel masks.....	92
12.2.2 Flow control	93
12.2.3 Isochronous data transformation.....	93
12.3 Error logs.....	94

Tables

Table 1 – Data transfer speeds.....	27
Table 2 – Management request functions.....	34
Table F-1 – SAM-2 Service responses	117

Figures

Figure 1 – Bit ordering within a byte	9
Figure 2 – Byte ordering within a quadlet	9
Figure 3 – Quadlet ordering within an octlet	9
Figure 4 – CSR specification example	10
Figure 5 – State machine example	12
Figure 6 – Linked list of ORB's	14
Figure 7 – Directly addressed data buffer.....	15
Figure 8 – Indirectly addressed data buffer (<i>via</i> page table)	15
Figure 9 – Components of an isochronous stream (direct-access target).....	18
Figure 10 – Stream engine block diagram.....	19
Figure 11 – Address pointer.....	23
Figure 12 – ORB pointer.....	23
Figure 13 – ORB family tree	24
Figure 14 – ORB format	24
Figure 15 – Dummy ORB	25
Figure 16 – Normal command block ORB.....	26
Figure 17 – Stream command block ORB.....	27
Figure 18 – Stream control ORB	29
Figure 19 – Channel mask	30
Figure 20 – Channel configuration map entry.....	31
Figure 21 – Management ORB.....	33
Figure 22 – Login ORB	35
Figure 23 – Login response	36
Figure 24 – Query logins ORB.....	36
Figure 25 – Query logins response format	37
Figure 26 – Create stream ORB.....	38
Figure 27 – Create stream response.....	39

Figure 28 – Reconnect ORB.....	40
Figure 29 – Logout ORB.....	41
Figure 30 – Task management ORB.....	41
Figure 31 – Page table element (unrestricted page table).....	42
Figure 32 – Page table element (when <i>page_size</i> equals four)	43
Figure 33 – Status block format.....	44
Figure 34 – MANAGEMENT_AGENT format	50
Figure 35 – AGENT_STATE format.....	52
Figure 36 – AGENT_RESET format.....	52
Figure 37 – ORB_POINTER format	53
Figure 38 – DOORBELL format.....	54
Figure 39 – UNSOLICITED_STATUS_ENABLE format	54
Figure 40 – Configuration ROM hierarchy	55
Figure 41 – Bus information block format.....	56
Figure 42 – Module_Vendor_ID entry format.....	57
Figure 43 – Node_Capabilities entry format	57
Figure 44 – Unit_Directory entry format.....	58
Figure 45 – Unit_Spec_ID entry format	58
Figure 46 – Unit_SW_Version entry format	58
Figure 47 – Command_Set_Spec_ID entry format.....	59
Figure 48 – Command_Set entry format	59
Figure 49 – Command_Set_Revision entry format	59
Figure 50 – Firmware_Revision entry format.....	60
Figure 51 – Management_Agent entry format.....	60
Figure 52 – Logical_Unit_Characteristics entry format.....	60
Figure 53 – Logical_Unit_Directory entry format	61
Figure 54 – Logical_Unit_Number entry format.....	61
Figure 55 – Unit_Unique_ID entry format	62
Figure 56 – Unit unique ID leaf format.....	63
Figure 57 – Fetch agent initialization with a dummy ORB	70
Figure 58 – Fetch agent state machine	72
Figure 59 – CYCLE MARK format	85
Figure 60 – Format for recorded isochronous data	86
Figure 61 – NULL packet format.....	87
Figure 62 – Cycle mark index format.....	87
Figure 63 – Common isochronous packet (CIP) format	88
Figure 64 – Two-quadlet CIP header format.....	88
Figure 65 – Source packet header format	89
Figure 66 – Synchronization time (<i>syt</i>) format	89
Figure 67 – Unsolicited status format for isochronous errors	95
Figure B-1 – SCSI command block ORB	99
Figure B-2 – SCSI control byte	99
Figure B-3 – Status block format for SCSI sense data	100
Figure C-1 – Set password ORB	106
Figure D-1 – Basic configuration ROM	109
Figure D-2 – SCSI configuration ROM	111
Figure G-1 – AV/C command frame ORB	123
Figure G-2 – Status block format for AV/C response frame	123
Figure G-3 – AV/C unit directory.....	124

Annexes

Annex A (normative) Minimum Serial Bus node capabilities	97
Annex B (normative) SCSI command and status encapsulation.....	99
Annex C (normative) Security extensions.....	105
Annex D (informative) Sample configuration ROM.....	109
Annex E (informative) Serial Bus transaction error recovery.....	113
Annex F (informative) SCSI Architecture Model.....	115
Annex G (informative) AV/C Encapsulation	122

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

This standard defines such a transport protocol within the domain of Serial Bus, IEEE Std 1394-1995, that is designed to permit efficient, peer-to-peer operation of input output devices (disks, tapes, printers, *etc.*) by initiator(s) such as operating systems or embedded applications. Vendors that wish to implement devices that connect to Serial Bus may follow the requirements of this and other standards to manufacture an SBP-2 compliant device.

This standard was developed by T10 during 1996 and 1997. Although some early SBP mock-up devices were demonstrated in 1993, significant proof-of-concept, in the form of prototype implementations, has proceeded contemporaneously with the development of this standard.

There are seven annexes in this standard. Annexes A, B and C are normative and part of this standard. Annexes D, E, F and G are informative and are not considered part of this 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 I Street NW, Suite 200, Washington, DC 20005-3922.

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

James D. Converse, Chair
Donald C. Loughry, Vice-chair
Joanne M. Flanagan, Secretary

<i>Organization Represented</i>	<i>Name of Representative</i>
American Nuclear Society	Geraldine C. Main
AMP, Inc.	Edward Kelly
Apple Computer	Karen Higginbottom
Association of the Institute for Certification of Professionals.....	Kenneth Zemrowski
AT&T/NCR	Thomas W. Kern
Boeing Company.....	Catherine Howells
Bull HN Information Systems, Inc.....	William George
Compaq Computer Corporation	James Barnes
Digital Equipment Corporation	Delbert Shoemaker
Eastman Kodak	James D. Converse
GUIDE International	Frank Kirshenbaum
Hewlett-Packard	Donald C. Loughry
Hitachi America, Ltd.	John Neumann
Hughes Aircraft Company.....	Harold L. Zebrack
IBM Corporation	Joel Urman
National Communication Systems	Dennis Bodson
National Institute of Standards and Technology	Robert E. Roundtree
Northern Telecom, Inc.	Mel Woinsky
Neville & Associates	Carlton Neville
Recognition Technology Users Association.....	Herbert P. Schantz
Share, Inc.	Gary Ainsworth
Sony Corporation	Michael Deese
Storage Technology Corporation	Joseph S. Zajackowski
Sun Microsystems	Scott Jameson

3M Company	Eddie T. Morioka
Unisys Corporation	John L. Hill
US Department of Defense.....	William C. Rinehuls
US Department of Energy.....	Alton Cox
US General Services Administration.....	Douglas Arai
Wintergreen Information Services	Joun Wheeler
Xerox Corporation.....	Dwight McBain

Technical Committee T10 on Lower Level Interfaces, which developed and reviewed this standard, had the following members:

John B. Lohmeyer, Chair	I. D. Allan	G. Milligan
Lawrence J. Lamers, Vice-chair	P. D. Aloisi	C. Monia
Ralph O. Weber, Secretary	R. Bellino	D. Moore
	C. Brill	I. Morrell
	R. Cummings	N. Nadershahi
	Z. Daggett	C. Nieves
	J. Dambach	E. Oetting
	R. Freeze	D. Pak
	E. A. Gardner	K. W. Parker
	D. Guss	G. Penokie
	K. J. Hallam	D. Piper
	E. Haske	G. Porter
	P. Jadeja	R. Reisch
	P. Johansson	J. R. Sims
	G. Johnsen	R. N. Snively
	S. Jones	G. R. Stephens
	C. Kephart	C. Tashbook
	M. Kuhlmeier	P. Tobias
	T. J. Kulesza	T. Totani
	A. Littlewood	D. Wagner
	B. Masterson	R. Wagner
	W. P. McFerrin	D. Wallace
	J. McGrath	J. L. Williams
	P. McLean	M. Wingard
	G. McSorley	K. Wolfgang
	P. Mercer	A. Yang

Revision history

Revision 1 (July 17, 1996)

First release of working draft.

Revision 1a (August 13, 1996)

Changes were incorporated from *ad hoc* discussions with diverse participants. These were presented at the Redmond, WA, SBP-2 working group meeting for discussion.

Data structure locations have been constrained to enable cost-reductions in target silicon. ORB's and associated parameter and response buffers shall be in the same node as the initiator that logged-in to the target. The same restriction shall also apply to the status FIFO. In a similar fashion, the data buffer and the page table that describe it shall reside both in the same node—although this node does not have to be the same as the initiator's.

The ORB fields that describe the data buffer and page table were enhanced to permit the description of data transfer alignment requirements in the case where the data buffer is directly addressable as a contiguous range of Serial Bus addresses.

Interrupt notification was modified to permit the return of status to be optional. An error condition overrides this parameter; a status block shall always be stored in the event of an error.

The ORB data structures were modified to compact the stream CDB and stream control ORB's to 32 bytes from 64 bytes.

New management ORB's have been defined for security management and access control. The accompanying work in section 8 still remains to be completed.

The status block has been expanded to permit the return of autosense data when appropriate to the device class. Targets are permitted to return portions of the status block when appropriate.

Unsolicited status was added as a feature, along with an interlock through a new register, STATUS_ACKNOWLEDGE, to let the initiator pace the receipt of unsolicited status reports.

The login and management agents have been collapsed into one agent, the management agent. Requests are signaled to the management agent *via* a new register, the MANAGEMENT_AGENT register, whose address is obtained from configuration ROM.

The target fetch agent (for normal CDB, stream CDB and stream control requests) has been enhanced to permit its reactivation from a SUSPENDED state by a single write to the ORB_POINTER register. This is an improvement over the previous approach where a write to the DOORBELL register would cause the fetch agent to refetch the *next_ORB* address before fetching the new request. A

Serial Bus transaction is eliminated and the restart latency is significantly improved.

A new configuration ROM entry, Unit_Unique_ID was defined to support SBP-2 devices that have multiple Serial Bus connections.

The basic task management model, discussed by the T10 SCSI-3 working group in Colorado Springs, CO, July 17, 1996, is now part of the draft.

Revision 1b (September 9, 1996)

Changes were incorporated as a result of working group discussions in Redmond, WA and were subsequently presented in Natick, MA.

The definition of a logical unit has been expanded. A target shall always implement logical unit zero.

Alignment restrictions on SBP-2 data structures (*i.e.*, anything referenced by an address pointer in an ORB other than the data buffer itself) have been relaxed from 16- to 4-byte alignment.

The normal and stream CDB ORB's have been simplified to permit a variable length CDB to follow the first five quadlets of the ORB, whose definition remains constant. This eliminated the need for both a 32- and 64-byte ORB.

The names *page_table* and *page_table_elements* have been changed throughout the document to *data_descriptor* and *data_size*, respectively. The meaning and usage of these fields has not changed. These global changes are not marked with change bars.

Block read transactions used to access page tables shall not cross page alignment boundaries expressed as 2^{page_size+8} bytes.

SBP-2 status has been redefined into two parts, one dependent upon the command set of the device and another used to present transport protocol status common to all devices.

As a result of discussions in Redmond, the fetch agent CSR's have been simplified and mistakes in the fetch agent state machine corrected. The figure that illustrates the fetch agent state machine and the accompanying text have been relocated to be closer to the descriptions of the usage of target fetch agents by initiators.

Section 8 has been expanded to document the usage of the login and security ORB's defined in 5.1.4.

Revision 1c (September 18, 1996)

Changes made *per* working group discussions in Natick, MA.

The acronym CDB has been changed throughout to command block or *command_block*, as appropriate. This global change is not marked by change bars.

The conformance glossary has been expanded to define the terms “reserved” and “ignored” and to clarify the implications of “shall.”

A note has been added to emphasize that device designers are encouraged to use 32-byte ORB's.

The *fetchable* bit in the status block has been renamed *end_of_list* and its meaning has been redefined. The status block has been modified to permit SBP-2 errors to be reported concurrently with command set errors. The *sense_key*, *asc* and *ascq* fields have been deleted and redefined as command set-dependent.

The MANAGEMENT_AGENT register has been redefined so that a write transaction, rather than a lock, is used to signal a request to the target.

The *doorbell*, *fetched* and *status_acknowledge* bits have been removed from AGENT_STATE register.

Configuration ROM definitions in the unit directory have been modified and a logical unit directory added to permit greater flexibility in the specification of targets that implement multiple logical units. The sample configuration ROM in the informative annex reflects the changes.

A new clause has been added to section 9 to describe the expected use of a target fetch agent by the BIOS or similar single-threaded application at an initiator.

The section on task management has been updated to improve clarity and to indicate that support for task management ORB's with a *function* of ABORT TASK is optional. Targets are still required to recognize an abort task request when the initiator sets the value of *rq_fmt* to three.

A normative annex has been added to specify the minimum Serial Bus requirements for both initiators and targets.

Revision 1d (October 5, 1996)

Editorial comments discussed in Irvine, CA, have been incorporated in this revision.

The definition of “reserved” has been changed so that a target shall not check the values of reserved fields.

The *notify* bit is advisory. That is, a target may return status even if *notify* is zero.

The initiator shall insure that *max_payload* does not specify a maximum data transfer larger than the speed code permits.

The circumstances under which a target may retry a block write transaction to an initiator's *status_FIFO* have been clarified.

The requirement for targets to implement the STATE_CLEAR.*dreq* bit has been stated in Annex A.

Revision 1e (November 9, 1996)

Minor editorial changes throughout, *per* discussions in Redmond, WA. The name of the CURRENT_ORB register has been changed to ORB_POINTER; this global change is unmarked by change bars.

After substantial discussion, the *ad hoc* working group concluded that security issues are best handled at the command set level. SBP-2 need provide only an access control mechanism that is sufficient to validate the actual identity of the initiator, EUI-64, and to provide a means whereby an initiator that had access rights before a Serial Bus reset has priority to reestablish the same access rights ahead of other, competing initiators. As a consequence, section 8 has been substantially revised and corresponding changes made to the data structure and configuration ROM descriptions in 5.1.4 and 7.4.8 respectively.

The fetch agent state machine diagram has been updated to simplify the actions a target shall take upon a write to the DOORBELL register.

The error conditions under which a target shall not attempt a retry of a block write transaction to store completion status have been clarified.

In Annex A, the target is required to support 8-byte block read and block write requests only for the MANAGEMENT_AGENT or ORB_POINTER register.

Former Annex B, "SCSI-3 Architecture Model compliance," has been removed to a separate document under development by T10 that includes a description of the use of SBP-2 facilities to implement SCSI devices.

Revision 1f (November 14, 1996)

The definition of "reserved" has been updated to bring it into conformance with contemporaneous standards such as T10 Project 1048D, SCSI-3 Multimedia Commands.

A *function* value for management ORB's has been set aside for command set-dependent use.

New *password* and *password_length* fields have been defined in the login ORB. The usage of these fields is command set-dependent but is intended to permit additional validation of the login ORB by a target.

An *exclusive* bit has been defined in the login ORB. When *exclusive* is set to one it causes multiple initiator targets to behave as if they supported only one login at a time.

In order to enable lower cost target hardware implementations, the format of the page table has been expanded and redundant information has been added. The net result is that the parsing of page tables may be normalized by target hardware. A requirement for octlet alignment of the page table elements was also added.

Targets shall not support broadcast write requests except as already required by IEEE Std 1394-1995 or future standards.

Logout requests are to be rejected if the *source_ID* does not match that of the currently logged-in initiator.

Revision 1g (December 4, 1996)

The response status returned by a target when an ORB with *rq_fmt* equal to three is processed (also known as a dummy ORB) is REQUEST ABORTED.

The previous revision had errors in the description of constraints that apply to page table elements, dependent upon their position within the page table. These errors have been corrected.

The descriptions of login and logout in section 8 have been clarified.

Portions of 10.4.1 have been rewritten in a simpler fashion that also permits greater target implementation flexibility in response to a task management ORB with the ABORT TASK *function*.

Revision 2 (January 9, 1997)

Subsequent to a vote by the T10 plenary to stabilize portions of SBP-2, this revision has been prepared; it is essentially identical to Revision 1g but without the change bars.

The sections stabilized by the plenary *exclude* the portions of SBP-2 concerned with isochronous data streams. The stabilized sections are enumerated below:

Section	Description
1	Scope and purpose
2	Normative references
3	Definitions and notation
4	Model (with the exception of 4.5)
5	Data structures (with the exception of 5.1.2.2 and 5.1.3)
6	Control and status registers (with the exception of 6.5)
7	Configuration ROM
8	Access (with the exception of 8.2.2)
9	Command execution
10	Task management
Annex A	Minimum Serial Bus node capabilities
Annex B	Sample configuration ROM

For readers unfamiliar with T10 process, stabilization is a significant milestone in the development of a standard. Once a document or portions thereof are stabilized they are not to be modified unless either a) there is a demonstrable flaw in the draft standard or b) the changes are agreed to by a two-thirds vote of the T10 plenary in which at least half of the membership votes.

Revision 2a (February 10, 1997)

Accredited Standards Committee X3 – Information Technology has changed its name to the National Committee for Information Technology Standardization (NCITS). The global changes are not marked by change bars.

Various typographical errors noted by reviewers have been corrected and marked with change bars.

The model of a target has been enhanced by a new clause that describes both direct and indirect access to the data buffer and by new material in the isochronous sections. As part of this, the name “isochronous login” has been changed to “create stream.”

The MANAGEMENT_AGENT and ORB_POINTER registers are required to support both block reads and block writes.

The numeric values used to encode data type in the common format for recorded isochronous data have been changed to align them with *tcode* values used by Serial Bus.

The stream control function, SET CHANNEL MASK, is restricted to only those times when an isochronous stream is paused or stopped.

Error reporting for isochronous streams has been simplified. Instead of an error log accumulated in initiator memory, unsolicited status is stored each time an error occurs.

Annex B (informative), “Status block for sense data”, has been added. Although this annex is informative within the context of SBP-2 its purpose is to provide material that may be normatively referenced from other command set standards, such as ATAPI, MMC-2 or SCSI.

Revision 2b (February 27, 1997)

The plug control registers (PCR's) and the procedures for their use have been removed because they do not add to the functionality already provided by SBP-2 stream control requests.

The model of the stream engine has been enhanced to show a functional component that filters channels from previously recorded isochronous data before it is made available to the Serial Bus LINK.

Fields in the stream command block ORB have been modified to permit cycle mark synchronization locations to be specified. This has utility when an SBP-2 device has lost current position information within an isochronous stream as a result of command block error(s).

Channel configuration information is conveyed as a single, 64-entry channel map instead of the previous method that required a stream control ORB to configure each channel.

A new field, *delta_time*, has been added to the CREATE STREAM request. Upon playback, *delta_time* is used to shift the time-stamps of certain isochronous data formats into the future.

A note has been added to recommend that writes to the target's RESET_START register be qualified by the *source_ID* of the write in order to restrict access to currently logged-in initiators.

Behavioral differences between normal command block agents and the pairs of stream command block and stream control agents allocated by a create stream request have been clarified. This includes a redefinition of ORB_POINTER behavior such that its contents are unchanged by a bus reset. Stream task sets are not aborted by a bus reset—they continue to execute but retain completion status until a reconnect by the initiator.

Support for abort task by means of task management ORB's with a *function* of ABORT TASK is mandatory for stream command block and stream control fetch agents.

While recording isochronous data, the target shall detect missed isochronous cycle(s) and synthesize a CYCLE MARK packet to record on the medium.

Errors in the description of the isochronous data interchange format that are not consonant with the CIP format specified by IEC 61883/FDIS have been corrected.

Revision 2c (March 26, 1997)

The name Command_Set_Version has been changed, globally, to Command_Set; the changes are not marked.

References to REQUEST ABORTED have been removed since this response code has never been defined. In conjunction, the description of abort task processing in 10.4.1 has been clarified to distinguish between dummy ORB completed and request aborted status.

The code values defined for the status block have been expanded to permit more accurate description of Serial Bus transport errors. Errors detected by a stream controller may also be reported by the use of unsolicited status.

A target's permission to store unsolicited status has been modified so that it is initially disabled at the time of a login or create stream request. Initiators that wish to accept unsolicited status may subsequently write to the UNSOLICITED_STATUS_ENABLE register each time a single instance of unsolicited status may be stored.

A description of the optional LOGICAL UNIT RESET function has been added to the task management section.

A tabular summary of fetch agent and task set states has been added to the section on task management. The table collects in one place target state information related to significant events such as power reset, bus reset or task management functions.

All of the above technical changes to stabilized portions of the draft standard were approved by the T10 plenary March 13, 1997, in San Diego, CA.

Revision 2d (June 10, 1997)

This draft reflects editorial changes proposed by both the *ad hoc* and editors' working groups in Natick, MA.

The model for isochronous targets is changed to make the stream control agent optional. Targets may provide the same facilities by other methods, such as plug control registers or commands served by the stream command block agent.

The relationship between status FIFO addresses explicitly specified by management ORB's and those implicitly associated with other ORB's has been clarified; there has been no technical change from the original intent of the working group.

The optional cycle mark index is described in the section on isochronous data interchange format.

The isochronous data transformation procedures specified in 12.2.3 have been corrected.

Revision 2e (July 23, 1997)

In addition to the technical changes described below, this draft incorporates editorial changes and typographical corrections from the *ad hoc* working group meeting in San Jose, CA.

The definition of the *data_descriptor* field in the normal command block ORB has been clarified to show that a directly addressed data buffer can start at any byte address.

Page table formats have been extended from the original (normalized) format to include a less restrictive format.

Additional information was added to the section on configuration ROM to clarify target behavior at the time of a power reset. References to the node unique ID leaf have been removed here and elsewhere in anticipation of P1394a changes that make this ROM information optional. A new entry, Firmware_Revision, has been added for use by manufacturers.

Target behavior after power reset is also specified more precisely in Annex A. Security requirements for the target have been added to this annex.

All of the above changes to stabilized portions of the draft standard were approved by the T10 plenary July 17, 1997, in Colorado Springs, CO.

The *idf* field was added to the CREATE STREAM request to control the format of recorded isochronous data.

Time stamp transformations for CIP format isochronous data were corrected to reflect the fact that the *cycle_offset* never requires transformation.

Annex D (informative) was added to create a location for advice to implementers about Serial Bus errors for which the error recovery strategy may not be obvious.

Revision 2f (September 7, 1997)

In addition to the technical changes described below, this draft incorporates editorial changes and typographical corrections from the *ad hoc* working group meeting in Longmont, CO.

The extent of configuration ROM covered by the CRC in the first quadlet is not specified by this standard; see instead IEEE Std 1394-1995.

Former Annex B (informative), "Status block for sense data", has been removed since it is substantially replicated by the new B.2 (see below).

New Annex B (normative) has been added to specify how devices that utilize the SCSI command set(s), status and sense data are to be implemented within SBP-2. Annex C (informative), which includes an example of SBP-2 configuration ROM, has been expanded to include examples specific to SCSI devices. Annex D (informative) explains how SCSI Architecture Model concepts are expressed within SBP-2.

New Annex E (informative) describes how AV/C command and response frames may be transported by SBP-2 instead of Function Control Protocol (FCP).

Revision 2g (September 15, 1997)

In addition to the technical changes described below, this draft incorporates editorial changes and typographical corrections from the editorial session and *ad hoc* working group meeting in Nashua, NH.

Contradictory descriptions of *ack_tardy* in 7.1 have been corrected.

The configuration ROM definition of *login_timeout* has been extended to all task management requests and the field renamed to *mgt_ORB_timeout*.

The above changes to stabilized portions of the draft standard were approved by the T10 plenary September 11, 1997, in Nashua, NH.

If unsolicited status is used to report the progress of commands, implementers are advised to pick reporting intervals that balance Serial Bus utilization against the usefulness of the information.

A definition of "working set," part of the task set kept in a target's local storage, has been added to the glossary.

Revision 2h (November 10, 1997)

A new annex, "Security extensions," provides a normative, but optional, method to add password protection to the SBP-2 access protocols. A change was also required to add a new management ORB function code in a stabilized portion of the draft standard. These changes were approved by the T10 plenary November 6, 1997, in Palm Springs, CA.

Editorial changes to sections 11 and 12 as a result of the SBP-2 *ad hoc* meeting in Palm Springs as well as to A.3, "Target security," to match the language in IEEE P1394a Draft 1.1.

Revision 3 (November 10, 1997)

This revision has been prepared for ballot by T10 to approve or disapprove its forwarding to NCITS for further standards processing; it is essentially identical to Revision 2h but without the change bars.

American National Standard for Information Systems –

Serial Bus Protocol 2 (SBP-2)

1 Scope and purpose

1.1 Scope

This standard defines a command and data transport protocol for High Performance Serial Bus, as specified by IEEE Std 1394-1995. The transport protocol, Serial Bus Protocol 2 or SBP-2, conforms to the requirements of the aforementioned standard as well as to ISO/IEC 13213:1994, Control and Status Register (CSR) Architecture for Microcomputer Buses, and permits the exchange of commands, data and status between initiators and targets connected to Serial Bus.

1.2 Purpose

Original development work for Serial Bus Protocol (SBP) was initiated out of a desire to adapt SCSI capabilities and facilities to a serial environment. Serial interconnects offer a migration path for SCSI into the future because they may be better suited to cost reduction and speed increases than the parallel interconnects first utilized by SCSI.

As development of the standard progressed, the working group recognized the solutions provided by SBP-2 were of general applicability to large classes of Serial Bus peripheral devices. With this in mind, the development work was redirected to provide mechanisms for the delivery of commands, data and status independent of the command set or device class of the peripheral. SBP-2 provides a generic framework that may be referenced by other documents or standards that address the unique requirements of a particular class of devices. The enhanced goals set for the design of SBP-2 are ranked below:

- The protocol should permit the encapsulation of commands, data and status from a diversity of command sets, legacy as well as future, in order to preserve the investment in an existing application and operating system software base;
- The protocol should enable the initiator to form an arbitrarily large set of tasks without consideration of implementation limits in the target;
- The protocol should allow the initiator to dynamically add tasks to this set while the target is active in execution of earlier tasks. The addition of new tasks should not interfere with the target's processing of tasks currently active;
- Although the protocol should enable varying levels of features and performance in target implementations, strong focus should be kept on a minimal set deemed adequate for entry-level environments;
- Within the constraints posed by the preceding goal, the hardware and software design of the initiator should not be unduly affected by variations in target capabilities;
- The protocol should take advantage of features of Serial Bus that offer improvement of functionality or permit the inclusion of new functionality. In particular, the isochronous features of Serial Bus are one of its key differentiators and should be fully supported by the protocol; and
- In order to promote the scalability of aggregate system performance, the protocol should distribute the DMA context from the initiator adapter to the target devices.

Although SBP-2 has been designed for Serial Bus as currently specified by IEEE Std 1394-1995, the Technical Committee anticipates that it will be appropriate for use with future extensions to Serial Bus as they are standardized.

2 Normative references

The standards named in this section contain provisions which, through reference in this text, constitute provisions of this American National Standard. At the time of publication, the editions indicated were valid. All standards are subject to revision; parties to agreements based on this American National Standard are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below.

Copies of the following documents can be obtained from ANSI:

Approved ANSI standards;

Approved and draft regional and international standards (ISO, IEC, CEN/CENELEC and ITUT); and

Approved and draft foreign standards (including BIS, JIS and DIN).

For further information, contact the ANSI Customer Service Department by telephone at (212) 642-4900, by FAX at (212) 302-1286 or *via* the world wide web at <http://www.ansi.org>.

Additional contact information for document availability is provided below as needed.

2.1 Approved references

The following approved ANSI, international and regional standards (ISO, IEC, CEN/CENELEC and ITUT) may be obtained from the international and regional organizations that control them.

IEEE Std 1394-1995, Standard for a High Performance Serial Bus

ISO/IEC 9899:1990, Programming Languages—C

ISO/IEC 13213:1994, Control and Status Register (CSR) Architecture for Microcomputer Buses

2.2 References under development

At the time of publication, the following referenced standards were still under development.

IEEE P1394a, Draft Standard for a High Performance Serial Bus (Supplement)

IEC 61883/FDIS, Digital Interface for Consumer Electronic AV Equipment

T10 Project 1157D, ANSI NCITS.xxx-199x, SCSI Architecture Model 2 (SAM-2)

T10 Project 1225D, ANSI NCITS.xxx-199x, SCSI Primary Commands 2 (SPC-2)

3 Definitions and notation

3.1 Definitions

3.1.1 Conformance

Several keywords are used to differentiate levels of requirements and optionality, as follows:

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

3.1.1.2 ignored: A keyword that describes bits, bytes, quadlets, octlets or fields whose values are not checked by the recipient.

3.1.1.3 may: A keyword that indicates flexibility of choice with no implied preference.

3.1.1.4 reserved: A keyword used to describe objects—bits, bytes, quadlets, octlets and fields—or the code values assigned to these objects in cases where either the object or the code value is set aside for future standardization. Usage and interpretation may be specified by future extensions to this or other standards. A reserved object shall be zeroed or, upon development of a future standard, set to a value specified by such a standard. The recipient of a reserved object shall not check its value. The recipient of a defined object shall check its value and reject reserved code values.

3.1.1.5 shall: A keyword that indicates a mandatory requirement. Designers are required to implement all such mandatory requirements to assure interoperability with other products conforming to this standard.

3.1.1.6 should: A keyword that denotes flexibility of choice with a strongly preferred alternative. Equivalent to the phrase “is recommended.”

3.1.2 Glossary

The following terms are used in this standard:

3.1.2.1 byte: Eight bits of data.

3.1.2.2 command block: Space reserved within an ORB to describe a command intended for a logical unit that controls device functions or the transfer of data to or from device medium. The format and meaning of command blocks are outside of the scope of SBP-2 and are command set- or device-dependent.

3.1.2.3 device server: A function component of a logical unit responsible to execute tasks initiated by command blocks that specify data transfer or other device operations.

3.1.2.4 initial node space: The 256 terabytes of Serial Bus address space that is available to each node. Addresses within initial node space are 48 bits and are based at zero. The initial node space includes initial memory space, private space, initial register space and initial units space. See either ISO/IEC 13213:1994 or IEEE Std 1394-1995 for more information on address spaces.

3.1.2.5 initial register space: A two kilobyte portion of initial node space with a base address of FFFF F000 0000₁₆. Core registers defined by ISO/IEC 13213:1994 are located within initial register space as are Serial Bus-dependent registers defined by IEEE Std 1394-1995.

3.1.2.6 initial units space: A portion of initial node space with a base address of FFFF F000 0400₁₆. This places initial units space adjacent to and above initial register space. The CSR's and other facilities defined by unit architectures are expected to lie within this space.

3.1.2.7 isochronous channel: A relationship, identified by a channel number, between a node that is the talker and one or more nodes that are listeners. One isochronous packet, identified by the channel number, may be sent by the talker during each isochronous cycle. Channel numbers are allocated cooperatively through isochronous resource management facilities.

3.1.2.8 isochronous cycle: An operating mode of Serial Bus that occurs, on average, every 125 microseconds. During an isochronous cycle, the bus is available to isochronous talkers, only. Cooperative allocation of isochronous bandwidth guarantees a bounded worst-case latency for isochronous data.

3.1.2.9 kilobyte: A quantity of data equal to 2^{10} bytes.

3.1.2.10 listener: A node that receives an isochronous packet for an isochronous channel during an isochronous cycle. There may be zero, one or more listeners for any given isochronous channel.

3.1.2.11 logical unit: The part of the unit architecture that is an instance of a device model, *e.g.*, mass storage, CD-ROM or printer. Targets implement one or more logical units; the device type of the logical units may differ.

3.1.2.12 login: The process by which an initiator obtains access to a set of target fetch agents. The target fetch agents and their control and status registers provide a mechanism for an initiator to signal ORB's to the target.

3.1.2.13 login ID: A value assigned by the target during a login or create stream process. The login ID establishes a relationship between an initiator and a task set. In the case of a create stream request, the login ID (or, synonymously, the stream ID) establishes a relationship between an initiator and an isochronous stream. Either kind of login ID is used to identify subsequent requests from an initiator; in some cases the login ID is not present in the operation request block and its value is implicit.

3.1.2.14 node ID: The 16-bit node identifier defined by IEEE Std 1394-1995 that is composed of a bus ID portion and a physical ID portion. The physical ID is uniquely assigned as a consequence of Serial Bus initialization.

3.1.2.15 normal command block: A command block whose ORB specifies a data buffer address for the transfer of data to or from the device medium. Sometimes referred to as an asynchronous command block, since asynchronous Serial Bus transfers are used to effect the data transfer.

3.1.2.16 octlet: Eight bytes, or 64 bits, of data.

3.1.2.17 operation request block: A data structure fetched from system memory by a target in order to execute the command encapsulated within it.

3.1.2.18 quadlet: Four bytes, or 32 bits, of data.

3.1.2.19 receive: When any form of this verb is used in the context of Serial Bus primary packets, it indicates that the packet is made available to the transaction or application layers, *i.e.*, layers above the link layer. Neither a packet repeated by the PHY nor a packet examined by the link is "received" by the node unless the preceding is also true.

3.1.2.20 register: A term used to describe quadlet aligned addresses that may be read or written by Serial Bus transactions. In the context of this standard, the use of the term register does not imply a specific hardware implementation. For example, in the case of split transactions that permit sufficient time

between the request and response subactions, the behavior of the register may be emulated by a processor.

3.1.2.21 split transaction: A transaction that consists of separate request and response subactions. Subactions are considered separate if ownership of the bus is relinquished between the two.

3.1.2.22 status block: A data structure written to system memory by a target when an operation request block has been completed.

3.1.2.23 stream: An object that represents a target's functions and resources necessary to transfer isochronous data from one or more Serial Bus channels to the device medium (the target is a listener) or to transfer isochronous data from the device medium to one or more Serial Bus channels (the target is a talker).

3.1.2.24 stream command block: A command block whose ORB does not specify a data buffer address. Data transfer to or from the device medium requires coordination with stream control ORB's that mediate data transfer from or to Serial Bus by means of isochronous transactions.

3.1.2.25 stream ID: A synonym for login ID when the login ID has been returned by the target in response to a create stream request.

3.1.2.26 system memory: The portions of any node's memory that are directly addressable by a Serial Bus address and which accepts, at a minimum, quadlet read and write access. Computers are the most common example of nodes that make system memory addressable from Serial Bus, but any node, including those usually thought of as peripheral devices, may have system memory.

3.1.2.27 talker: A node that transmits an isochronous packet for an isochronous channel during an isochronous cycle. There shall be no more than one talker for any given isochronous channel.

3.1.2.28 task: A task is an organizing concept that represents the work to be done by a target to carry out a command encapsulated by an ORB. In order to perform a task, a target maintains context information for the task, which includes (but is not limited to) the command, parameters such as data transfer addresses and lengths, completion status and ordering relationships to other tasks. A task has a lifetime, which commences when the task is signaled to the target, proceeds through a period of execution by the target and finishes when completion status is stored at the initiator. While a task is active, it makes use of both target resources and initiator resources.

3.1.2.29 task set: A group of tasks available for execution by a logical unit of a target. There may be dependencies between individual tasks within the task set specified by this standard.

3.1.2.30 terabyte: A quantity of data equal to 2^{40} bytes.

3.1.2.31 transaction: An exchange between a requester and a responder that consists of a request and a response subaction. The request subaction transmits a Serial Bus transaction such as quadlet read, block write or lock, from the requesting node to the node intended to respond. Some Serial Bus commands include data as well as transaction codes. The response subaction returns completion status and sometimes data from the responding node to the requesting node.

3.1.2.32 unit: A component of a Serial Bus node that provides processing, memory, I/O or some other functionality. Once the node is initialized, the unit provides a CSR interface that is typically accessed by device driver software at an initiator. A node may have multiple units, which normally operate independently of each other. Within this standard, a unit is equivalent to a target.

3.1.2.33 unit architecture: The specification of the interface to and the services provided by a unit implemented within a Serial Bus node. This standard is a unit architecture for SBP-2 targets.

3.1.2.34 unit attention: A state that a logical unit maintains while it has unsolicited status information to report to one or more logged-in initiators. A unit attention condition shall be created as described elsewhere in this standard or in the applicable command set- and device-dependent documents. A unit attention condition shall persist for a logged-in initiator until a) unsolicited status that reports the unit attention condition is successfully stored at the initiator or b) the initiator's login becomes invalid or is released. Logical units may queue unit attention conditions; after the first unit attention condition is cleared, another unit attention condition may exist.

3.1.2.35 working set: The part of a task set that has been fetched from the initiator by the target and is available to the target in its local storage.

3.1.3 Abbreviations

The following are abbreviations that are used in this standard:

CDB	Command descriptor block
CIP	Common isochronous packet format
CSR	Control and status register
CRC	Cyclical redundancy checksum
DVCR	Digital video cassette recorder
EUI-64	Extended Unique Identifier, 64-bits
LUN	Logical unit number
MPEG	Motion picture experts group
ORB	Operation request block
SAM-2	SCSI Architecture Model 2
SBP-2	Serial Bus Protocol 2 (this standard itself)
SPC-2	SCSI Primary Commands 2

3.2 Notation

The following conventions should be understood by the reader in order to comprehend this standard.

3.2.1 Numeric values

Decimal, hexadecimal and, occasionally, binary numbers are used within this standard. By editorial convention, decimal numbers are most frequently used to represent quantities or counts. Addresses are uniformly represented by hexadecimal numbers. Hexadecimal numbers are also used when the value represented has an underlying structure that is more apparent in a hexadecimal format than in a decimal format. Binary numbers are used infrequently and generally limited to the representation of bit patterns within a field.

Decimal numbers are represented by Arabic numerals without subscripts or by their English names. Hexadecimal numbers are represented by digits from the character set 0 – 9 and A – F followed by the subscript 16. Binary numbers are represented by digits from the character set 0 and 1 followed by the

subscript 2. For the sake of legibility, binary and hexadecimal numbers are separated into groups of four digits separated by spaces.

As an example, 42, $2A_{16}$ and $0010\ 1010_2$ all represent the same numeric value.

3.2.2 Bit, byte and quadlet ordering

SBP-2 is defined to use the facilities of Serial Bus, IEEE Std 1394-1995, and therefore uses the ordering conventions of Serial Bus in the representation of data structures. In order to promote interoperability with memory buses that may have different ordering conventions, this standard defines the order and significance of bits within bytes, bytes within quadlets and quadlets within octlets in terms of their relative position and not their physically addressed position.

Within a byte, the most significant bit, *msb*, is that which is transmitted first and the least significant bit, *lsb*, is that which is transmitted last on Serial Bus, as illustrated below. The significance of the interior bits uniformly decreases in progression from *msb* to *lsb*.




Figure 1 – Bit ordering within a byte

Within a quadlet, the most significant byte is that which is transmitted first and the least significant byte is that which is transmitted last on Serial Bus, as shown below.




Figure 2 – Byte ordering within a quadlet

Within an octlet, which is frequently used to contain 64-bit Serial Bus addresses, the most significant quadlet is that which is transmitted first and the least significant quadlet is that which is transmitted last on Serial Bus, as the figure below indicates.




Figure 3 – Quadlet ordering within an octlet

When block transfers take place that are not quadlet aligned or not an integral number of quadlets. No assumptions can be made about the ordering (significance within a quadlet) of bytes at the unaligned beginning or fractional quadlet end of such a block transfer, unless an application has knowledge (outside of the scope of this standard) of the ordering conventions of the other bus.

3.2.3 Register specifications

This standard precisely defines the format and function of control and status registers, CSR's. Some of these registers are read-only, some are both readable and writable and some generate special side effects subsequent to a write.

In order to precisely define CSR's, their bit fields, their initial values and the effects of read, write or other transactions, the format illustrated by Figure 4 is used.

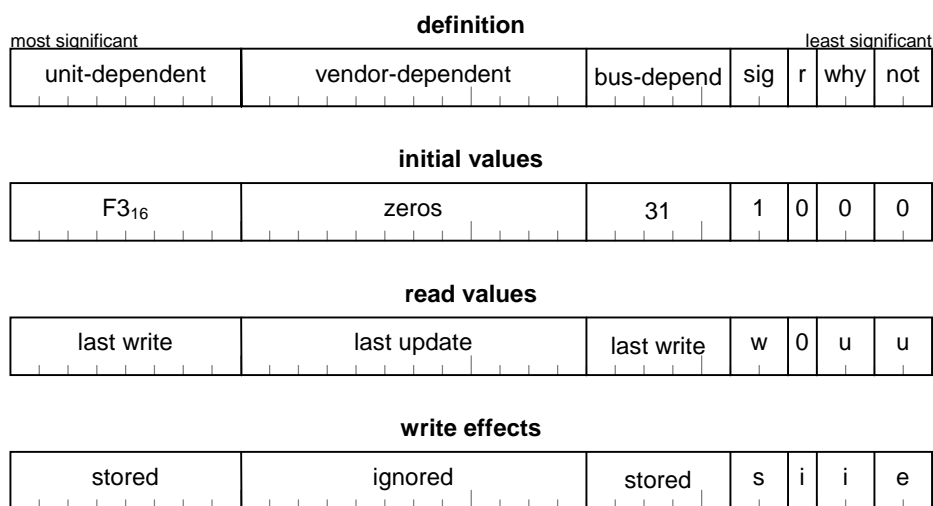


Figure 4 – CSR specification example

The register definition contains the names of register fields. The names are intended to be descriptive, but the fields are defined in the text; their function should not be inferred solely from their names. However, the following field names have defined meanings.

Name	Abbreviation	Definition
bus-dependent	bus-depend	The meaning of the field shall be defined by the bus standard, in this case IEEE Std 1394-1995
reserved	r	The field is reserved for future standardization (see definitions)
unit-dependent	unit-depend	The meaning of the field shall be defined by the organization responsible for the unit architecture
vendor-dependent	vendor-depend or v	The meaning of the field shall be defined by the node's vendor

CSR's shall assume initial values upon the restoration of power (a power reset) or upon a write to the node's RESET_START register (a command reset). If the power reset values differ from the command reset values, they are separately and explicitly defined. Initial values for register fields may be described as numeric constants or with one of the terms defined for the register definition. Values for register fields subsequent to a reset may be described in the same terms or as defined below.

Name	Abbreviation	Definition
unchanged	x	The field retains whatever value it had just prior to the power reset, bus reset or command reset.

In addition to numeric values for constant fields, the read values returned in response to a quadlet read transaction may be specified by the terms below.

Name	Abbreviation	Definition
last write	w	The value of the field shall be either the initial value or, if a write or lock transaction addressed to the register has successfully completed, the value most recently stored in the field. ¹
last update	u	The value of the field shall be that most recently updated by the node hardware or software. An updated field value may be the result of a write effect to the same register address, a different register address or some other change of condition within the node.

The effects of data written to the register shall be specified by the terms below.

Name	Abbreviation	Definition
effect	e	The value of the data written to the field may have an effect on the node's state, but the effect may not be immediately visible by a read of the same register. The effect may be visible in another register or may not be visible at all.
ignored	i	The value of the data written to the field shall be ignored; it shall have no effect on the node's state.
stored	s	The value of the data written to the field shall be immediately visible by a read of the same register; it may also have other effects on the node's state.

Reserved fields within a register shall be explicitly described with respect to initial values, read values and write effects. Initial values and read values shall be zero while write effects shall be ignored. CSR's that are not implemented, either because they are optional or they fall within a reserved address space, shall abide by these same conventions if a successful completion response is returned for a read, write or lock transaction.

3.2.4 State machines

All state machines in this standard are defined in the style illustrated by Figure 5.

¹ For clarity, read values for a field in a register that accepts lock transactions may be described as *last successful lock* rather than *last write*. However, the abbreviation in both cases remains *w*. Similar liberties may be taken with the use of *conditionally stored* in place of *stored* when the action occurs as the result of a lock transaction, but the corresponding one-letter abbreviation, *s*, is also unchanged.

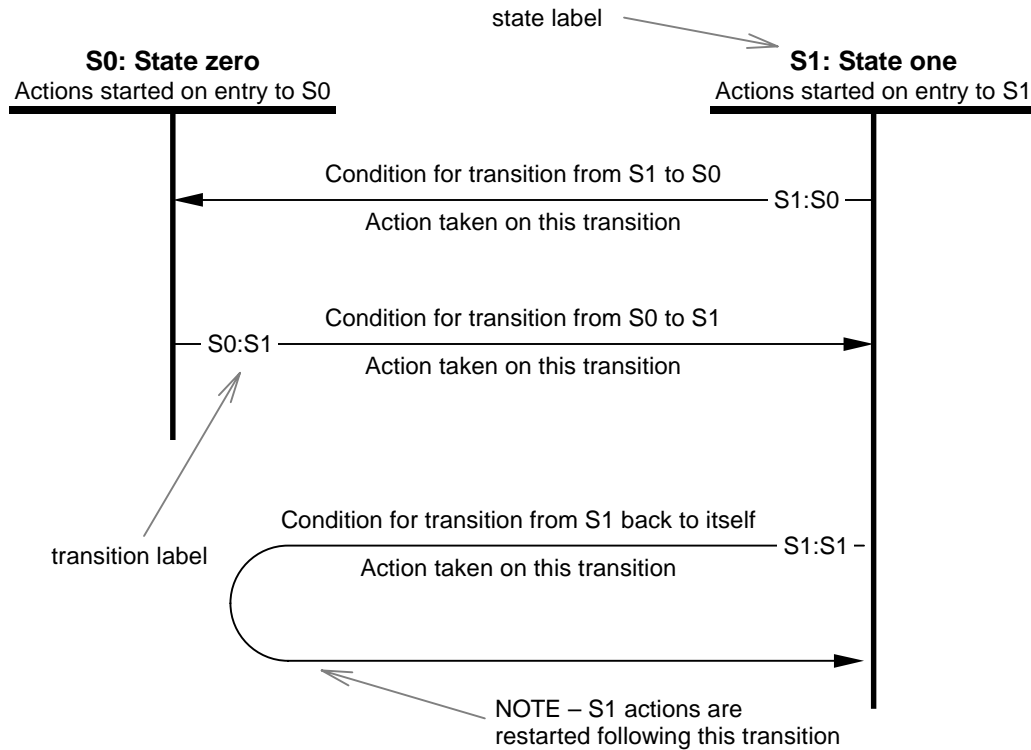


Figure 5 – State machine example

The state machines in this standard make three assumptions:

- Time elapses only within a discrete state;
- State transitions are conceptually instantaneous; the only actions taken during the transition are the setting of flags or variables and the sending of signals; and
- Each time a state is entered (or reentered from itself), the actions of that state are performed.

Multiple transitions may connect two states. In this case, the transitions are uniquely labeled by appending a character to the transition label, *e.g.*, S0:S1a and S0:S1b.

4 Model (informative)

Serial Bus Protocol 2 (SBP-2) is a transport protocol defined for IEEE Std 1394-1995, Standard for a High Performance Serial Bus. It defines facilities for requests (commands) originated by Serial Bus devices (initiators) to be communicated to other Serial Bus devices (targets) as well as the facilities required for the transfer of data or status associated with the commands. An SBP-2 device may assume the roles of initiator or target, either simultaneously or in succession. Commands and status may be transferred between the initiator and the target; data moves between the target and another device, which may be neither the initiator nor an SBP-2 device.

This clause is informative and describes components of the SBP-2 model. It is intended to enhance the usefulness of the other, normative parts of this standard. In addition to the information in this clause, users of this standard should also be familiar with the CSR architecture and Serial Bus standards.

4.1 Unit architecture

In CSR architecture and Serial Bus terminology, devices implemented to this standard (targets) are units. A Serial Bus node that implements a target has a unit directory in configuration ROM that identifies the presence and capabilities of the target.

The unit directory in configuration ROM permits initiators to detect the presence of targets during Serial Bus configuration, whether part of system initialization or subsequent to a Serial Bus reset. The node's 64-bit identifier, EUI-64, permits detected targets to be uniquely recognized despite changes in physical addresses that may occur as the result of Serial Bus resets.

4.2 Logical units

A logical unit is part of the unit architecture and is an instance of a device model, *e.g.*, mass storage, CD-ROM or printer. A logical unit consists of one or more device server(s) responsible to execute control or data transfer commands, zero or more stream controllers, one or more task sets that hold commands available for execution by the device server(s) or stream controller(s) and a logical unit number that is unique within the domain of the target.

Targets implement at least one logical unit, addressable as logical unit number (LUN) zero. Additional logical units may be implemented, which may be addressable by their logical unit numbers. The logical units may implement different device models; for example, a single unit architecture might contain both a CD-ROM logical unit and an associated medium-changer logical unit. The logical unit(s) are visible to the initiator, either as described by configuration ROM or as discoverable by command set-dependent requests directed to the target.

4.3 Requests and responses

Target actions, such as a disk read that transfers data from device medium to system memory, are specified by means of requests created by the initiator and signaled to the target. The request is contained within a data structure called an operation request block (ORB). The eventual completion status of a request is indicated by means of a status block stored by the target at an address provided by the initiator.

This standard defines several different formats for request blocks, whose principal uses are:

- to obtain access to target resources (login requests);
- to transport command blocks (normal and stream command block requests);
- to manage task sets or to release target resources (management requests); or
- to control the flow of isochronous data (stream control requests).

Login and management requests are directed to agents that can service only a single request at a time. The ORB's for the other requests, normal command block, stream command block and stream control, provide a field that contains the address of another ORB or a null pointer. This permits these requests to be in a linked list, as illustrated below.

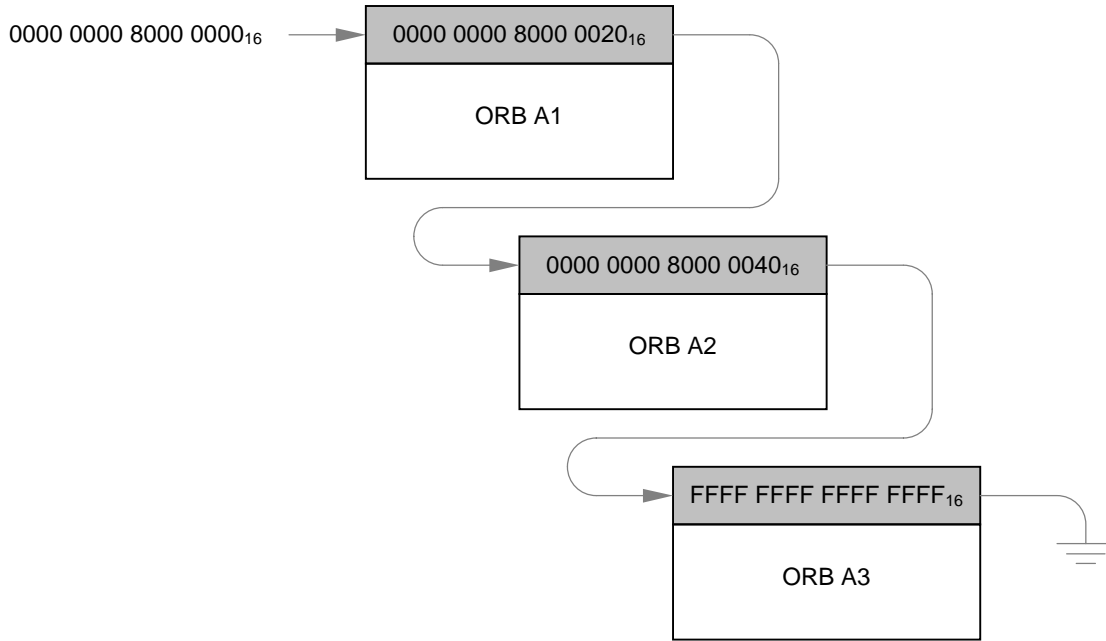


Figure 6 – Linked list of ORB's

Requests in a linked list are serviced by a target fetch agent, which reads the request(s) from initiator memory when the initiator signals the availability of request(s). The target may read ahead in the linked list; consequently the device server may reorder the execution of requests to improve performance.

When the request is completed, either in success or failure, the target stores a status block at an address specified by the initiator.

4.4 Data buffers

The ORB's described in the preceding clause contain the device command and the address of the data buffer for the command. The data buffer may be a single, contiguous buffer that is addressed directly by the ORB or it may be a collection of possibly disjoint segments that are addressed indirectly through a page table. The figures below illustrate both cases.

As an example, consider a command intended to transfer image data to a printer. If we assume that the image data is `308816` bytes long and the buffer starts at an address of `23 617416`, the relationship between the ORB and the data buffer might appear as in Figure 7.




Figure 7 – Directly addressed data buffer

In the preceding example, two fields in the ORB specify the 64-bit address of the data buffer and its length, in bytes. The data buffer is shown with a node ID of FFC0₁₆, which is node zero on the local bus. The printer uses block read transactions to fetch data from the buffer before printing; the maximum size of the data payload for each request is controlled by a field in the ORB. The dotted lines within the data buffer indicate page boundaries. Although the data buffer is contiguous, the printer is not permitted to cross a page boundary in any one block read request.

When the data buffer consists of disjoint segments, it is necessary to indirectly address the data buffer through a page table, as shown in Figure 8. This figure could be an illustration of data read from a disk into various pages of an initiator's file system cache. In the example, assume that 2960₁₆ bytes of data are to be read from disk.

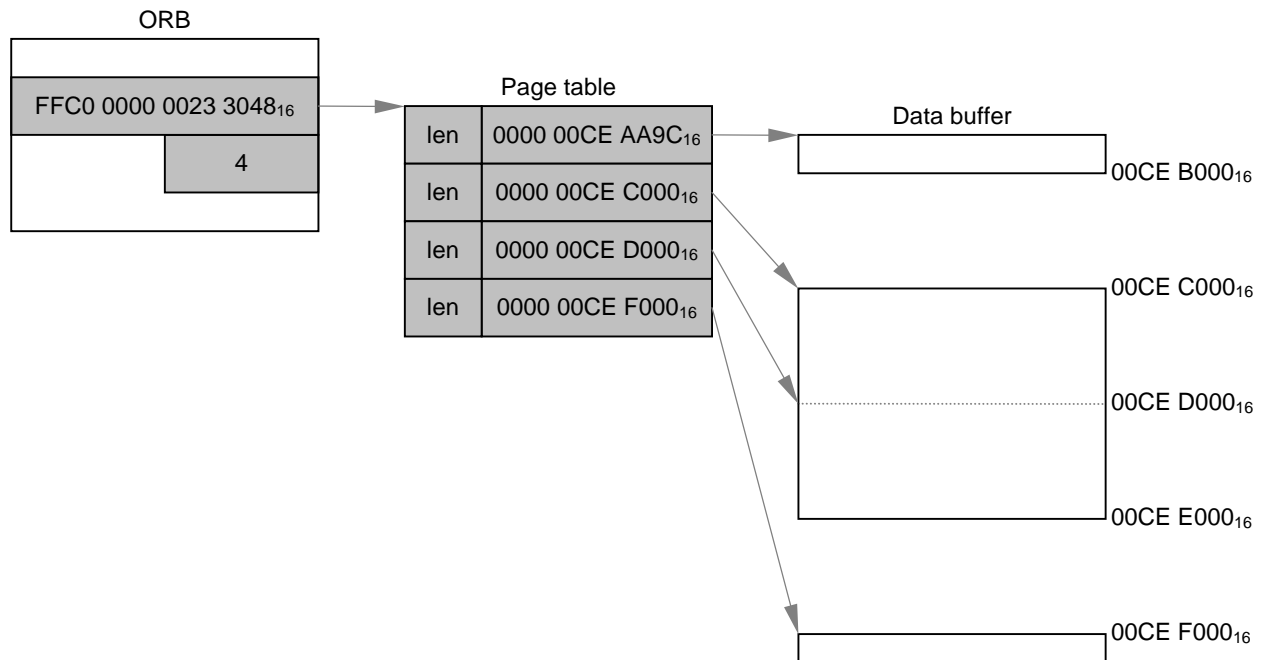


Figure 8 – Indirectly addressed data buffer (via page table)

The fields in the ORB that directly addressed a data buffer in the first example now point to a page table. Note that the ORB field that contains the data length when direct addressing is employed instead contains the number of elements in the page table—in this case, four. Each of the four page table elements points to the start of a segment of the data buffer. Each page table element also contains the length of the segment. The first segment ends on a page boundary, all other segments start on page boundaries (and the middle segments also end on page boundaries) while the last segment may end on any boundary. In this example, the segment lengths are 0564_{16} , 1000_{16} , 1000_{16} and $03FC_{16}$, respectively.

When a page table is used, both the page table and the data buffer it describes reside in the same node. The node ID of the page table, $FFC0_{16}$, is not repeated in the page table elements. The space that would have otherwise been occupied by the node ID instead is used to contain the length of each segment.

Another variant of page table format is permitted, called an unrestricted page table (or a scatter/gather list). In an unrestricted page table, data buffer segments may start on any boundary and may have arbitrary lengths: there is no underlying page size.

4.5 Target agents

A target agent is a facility that receives signals from the initiator that indicate the availability of requests. There are two types of target agent, one that can execute a single request at a time and the other that can manage queues (linked lists) of requests, as illustrated by Figure 6. In the first case, the initiator signals the request to the agent by means of a Serial Bus block write request with the address of the request. In the other case, the initiator appends new requests to an active list, rings a doorbell which causes the target agent to fetch the requests from system memory as target resources permit their execution.

Target agents that manage linked lists of requests utilize context maintained at both the initiator and target to fetch requests from memory. Once fetched, the request is locally available to the target for execution. The context consists of three elements:

- a linked list of ORB's at the initiator;
- a current ORB address at the target; and
- a doorbell at the target.

This standard defines procedures for both the initiator and the target that permits the addition of new requests to a linked list of ORB's while the target is actively fetching or executing previously enqueued requests. The procedures avoid the possibility of race conditions between the producer (initiator) and consumer (target) of the ORB's.

There are three defined target agents:

- management;
- command block; and
- stream control.

Management agents accept a variety of requests: login, create stream, task management and logout. Before making other requests, an initiator first completes a login *via* the management agent. Once this is done, the management agent will accept create stream requests and task management requests. The latter are directed to either a normal (asynchronous) task set or to a task set associated with an isochronous stream. Ultimately, management agents accept logout requests; these indicate the initiator's intent to release target resources previously acquired by a login or create stream request. Management agents service a single request at a time and do not support linked lists.

A successful login or create stream request returns the address of a command block agent. Command block agents service either normal or stream command block requests that are organized into linked lists. Each linked list is managed by a separate command block agent.

A create stream request may return the address of a stream control agent. This agent meters the flow of isochronous data to or from Serial Bus. Unless the target provides other facilities to meter this flow (*e.g.*, plug control registers as specified by IEEE P1394a), each stream requires both a stream command block agent and a stream control agent to coordinate operations. The time-critical nature of isochronous operations requires that stream control agents support linked lists of requests, just as command block agents.

4.6 Ordered and unordered execution

Targets may implement either an ordered or unordered model of task execution. The ordered model is usually appropriate for devices where the context of a command affects its execution, *i.e.*, the outcome of one command affects the subsequent command. A common example of a device with such command dependencies is a tape drive. The unordered model is usually appropriate for devices, such as mass storage, where no positional or other context information is inherited from one command to the next.

The ordered model specifies both that tasks shall be executed in order and that completion status shall be returned in the same order. A consequence of ordering that completion status for one task implicitly indicates successful completion status for all tasks that preceded it in the ordered list.

The unordered model permits the target to reorder active tasks without restriction. The actual execution sequence of tasks from any task set may bear no relationship to the order in which they were fetched. Unrestricted reordering places the responsibility for the assurance of data integrity on the initiator. If the integrity of data on the device medium could be compromised by unrestricted reordering involving a set of active tasks, $\{T_0, T_1, T_2, \dots T_N\}$ and a new task T' , the initiator shall wait until $\{T_0, T_1, T_2, \dots T_N\}$ have completed before appending T' to an active request list.

NOTE – In multitasking operating system environments, independent execution threads may generate tasks that have ordering constraints within each thread but not with respect to other threads. If this is the case, an initiator may manage the constraints of each thread yet still keep the target substantially busy. This avoids the undesirable latencies that occur if the target is allowed to become idle before new ORB's are signaled.

4.7 Streams

Streams are objects that are based upon the isochronous capabilities of Serial Bus. A stream consists of all of the target functions and resources that are necessary to transfer isochronous data from one or more Serial Bus channels to the device's medium (the target is a listener) or to transfer isochronous data from the device's medium to one or more Serial Bus channels (the target is a talker). The direction, listener or talker, of any stream is independent of any other stream. Within each stream all of the channels flow in the same direction.

Streams require Serial Bus resources as well as target resources. These include the aggregate bandwidth necessary for the stream, the channel numbers utilized by the stream and the isochronous connections that characterize the stream. An application allocates all necessary resources before activating a target isochronous stream.

A stream of isochronous data appears on Serial Bus as packet(s) during each isochronous cycle. This in turn is represented by an ordered byte stream of data on the device medium. The presentation of this data is controlled by ORB's that request data transfer to or from the medium and, optionally, other ORB's that control its flow on Serial Bus. The second type of ORB, the stream control ORB is not required if the target has other facilities to control the flow of isochronous data from or to Serial Bus.

The relationship between the different stream components during playback (for a target with direct-access capabilities) is given by Figure 9.

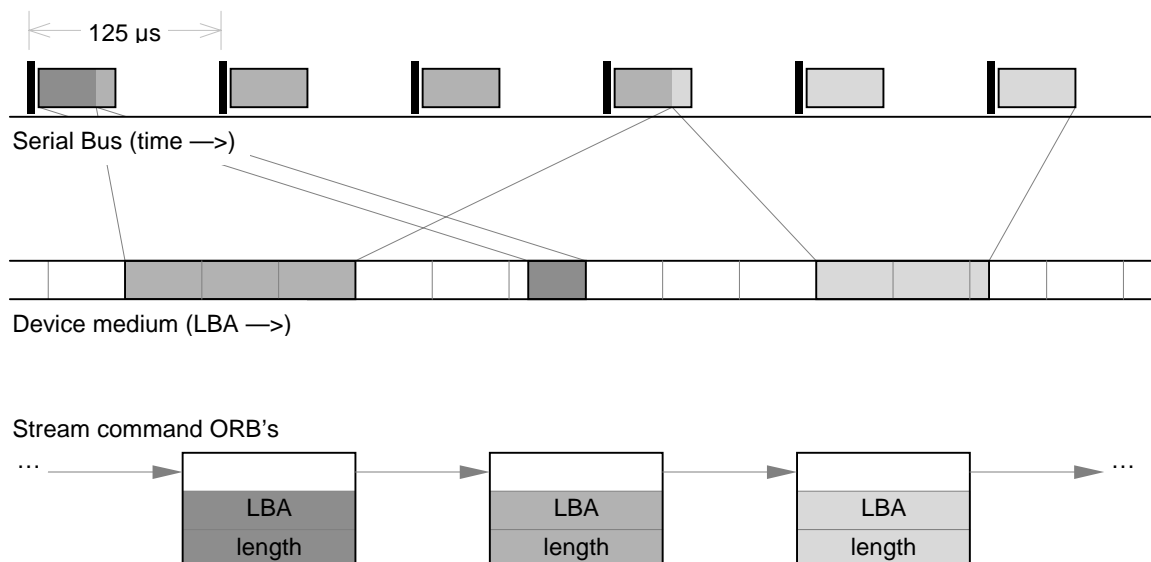


Figure 9 – Components of an isochronous stream (direct-access target)

The figure shows that the order of presentation of bytes in a stream is determined by the order of stream command ORB's—but that this order is independent of the location of the data on device medium. In this example, the size of the isochronous packet transmitted each cycle is determined by information previously recorded on the medium. The example shows a stream with only one channel (one packet per isochronous cycle) and a fixed packet size, but streams may consist of more than one channel and the packet size may be different each isochronous cycle.

Streams differ fundamentally from the data transfers described by normal command blocks in two important respects. First, streams do not require any address context for the transfer of data to or from system memory: stream data is identified by a channel number and the time-ordered location of the data within the stream. Second, streams permit flow control that is synchronized to time or other time-dependent events.

Because of these differences from normal operations, two functional components are required within the target to fully control a stream: a stream task set and a stream controller. The functions of a stream controller may be provided by a stream control agent or by methods outside the scope of this standard. Figure 10 illustrates the way in which these target components cooperate to form a functional unit, the stream engine.

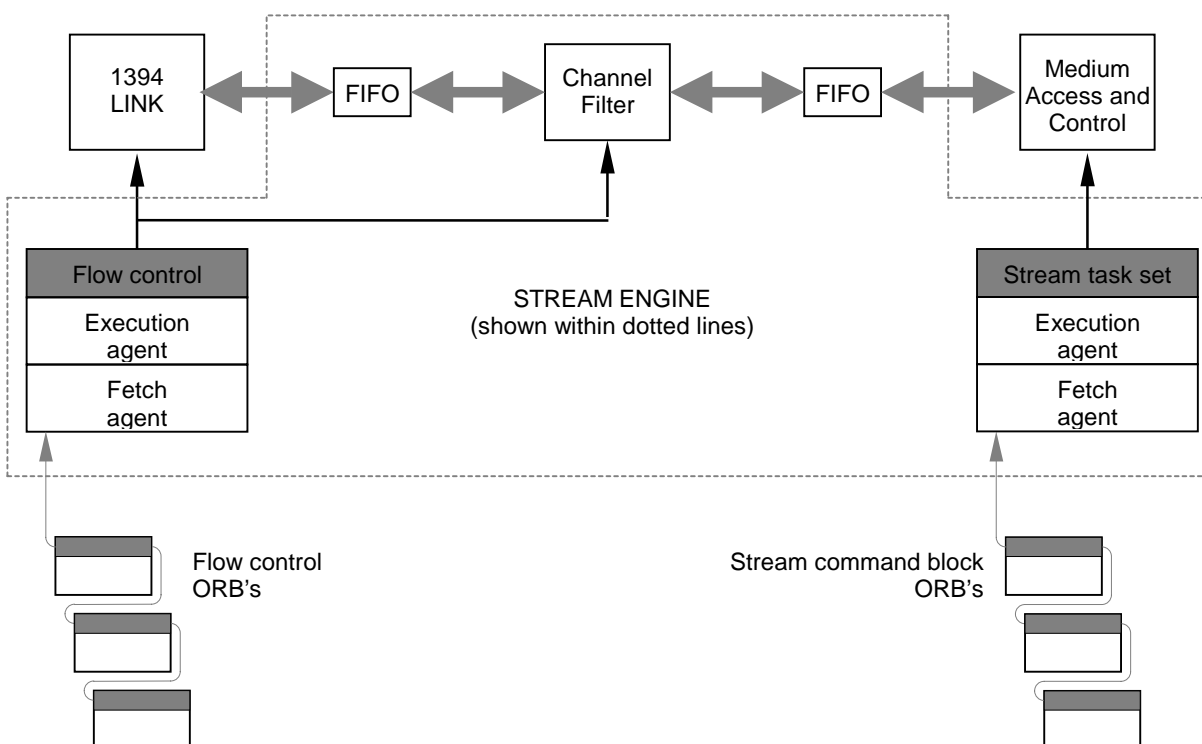


Figure 10 – Stream engine block diagram

The stream engine components that are shown within the dotted lines represent target resources necessary for each instance of an isochronous stream. The two execution agents each have separate queues of ORB's that direct their operation. The stream command block ORB's specify the length and location of isochronous data as it is transferred to or from the device medium while the flow control ORB's synchronize and meter the isochronous data on Serial Bus as the target listens or talks. These two components do not communicate directly with each other—they coordinate their operations by using a common isochronous data format for the information in the FIFO's that separate them. Although the details of the FIFO's (or buffers) are implementation-dependent, this model assumes that the byte ordering of the isochronous data is preserved.

NOTE – When a target does not implement a stream control agent, the left-hand side of the figure (which shows flow control ORB's and their fetch agent) is not present. Other facilities, such as plug control registers, provide a means to specify channel number(s) while the flow control functions START, PAUSE and STOP are typically part of the command set used in the stream command block ORB's.

4.7.1 Stream task set

Both a normal (or asynchronous) task set and a stream task set consist of a set of commands that request data transfer to or from a device's medium. A stream task set is different in two important respects:

- no system memory addresses; and
- implicit order relationships.

No system memory addresses are needed because data transferred to or from the device medium is associated with one or more Serial Bus isochronous channels. When isochronous data is read from the device medium it is made available, in order, to the stream controller described below. When isochronous

data is written to the device medium it is obtained, in order, from the stream controller. In neither case is a system memory address necessary.

Isochronous data is essentially time-ordered. As a consequence, the isochronous data transferred to or from the device medium must be presented in correct order. Therefore no reordering of isochronous commands is permitted within the task set associated with an isochronous stream and the failure of any one task requires that all subsequent tasks be aborted.

This requires not only that tasks in a stream task set are executed in order but also that their completion status reported in the same order.

4.7.2 Stream controller

By means that are implementation dependent, an ordered data pipe is assumed to exist between a target's stream task set and the associated stream controller. The function of the stream controller is to mediate the flow of isochronous data between this data pipe and Serial Bus.

The format of data transported through the data pipe is similar to that of Serial Bus isochronous packets. The data is identified by time stamps (cycle times) and channel numbers and the payload is described in terms of its length, in bytes.

The stream controller:

- filters isochronous data according to channel numbers;
- transforms time stamps and channel numbers in the isochronous data; and
- synchronizes the flow of the isochronous data with external, time-dependent events.

Any of these operations may take place whether the stream controller is a listener or a talker. Stream control ORB's that specify these operations are independent of the stream command block ORB's in the stream task set. The stream task set and the stream controller communicate with each other through the data pipe.

Stream controller actions may be queued by the target. This permits time-critical operations to be specified in advance and avoids latency problems that could arise if the stream controller could accept no more than one request at a time. Within the queue of requests to the stream controller, each is executed in order as the preceding stream control ORB completes.

4.7.3 Error reporting

In addition to the data transfer errors that may be encountered by any of the stream command block ORB's, errors may occur within the isochronous stream itself as it is transferred to or from Serial Bus. These errors might include, but are not limited to:

- a missing isochronous packet or cycle start indication;
- an isochronous packet with a header CRC error;
- when the target is a talker, an underflow in the availability of data from the stream command block ORB's that causes no data to be transmitted for one or more channels during an isochronous cycle; or
- when the target is a listener, an overflow in which isochronous data from Serial Bus must be discarded because of an internal buffer overflow or a lack of stream command block ORB(s) to transfer the data to the medium.

In any of these cases, the initiator may wish to ignore all errors, report all errors but continue the isochronous stream or report the first error and halt isochronous operations. The initiator selects the error reporting option, which determines whether or not the target may store an unsolicited status block at the time an error occurs.

5 Data structures

There are three classes of data structures defined by Serial Bus Protocol:

- operation request blocks (ORB's);
- page tables;
- status blocks.

These data structures may be allocated and initialized by an initiator in system memory at Serial Bus nodes. ORB's and status blocks shall be allocated at the initiator's node; page tables shall be allocated at the same node as the data buffer to which they refer.

All data structures defined by this standard shall be aligned on quadlet boundaries. These alignment requirements permit 64-bit address pointers within ORB's to conform to the format specified below.




Figure 11 – Address pointer

The *node_ID* field shall specify the Serial Bus node for which the address pointer is valid, as defined by IEEE Std 1394-1995. In many cases, additional constraints on the location of data structures render the information in *node_ID* redundant. In these cases, *node_ID* is considered a reserved field or is explicitly redefined for other uses.

The *offset_hi* and the *offset_lo* fields shall together specify the most significant 46 bits of the Serial Bus offset and shall be combined with two low-order bits of zero to derive the 48-bit Serial Bus offset.

ORB's shall be allocated at the initiator's node and may be organized into a linked list. Since the node ID is known for all ORB's in a given list, the address pointer format is redefined to reuse the *node_ID* field. An address pointer that references an ORB shall conform to the format below.




Figure 12 – ORB pointer

The *null* bit (abbreviated as *n* in the figure above) indicates a null pointer when it is one. In this case the target shall ignore the ORB offset fields.

5.1 Operation request blocks (ORB's)

All initiator requests for target actions are expressed within ORB's fetched by the target *via* Serial Bus read transaction(s). ORB formats vary according to use and may be viewed in hierarchical relationship to each other, as illustrated below.

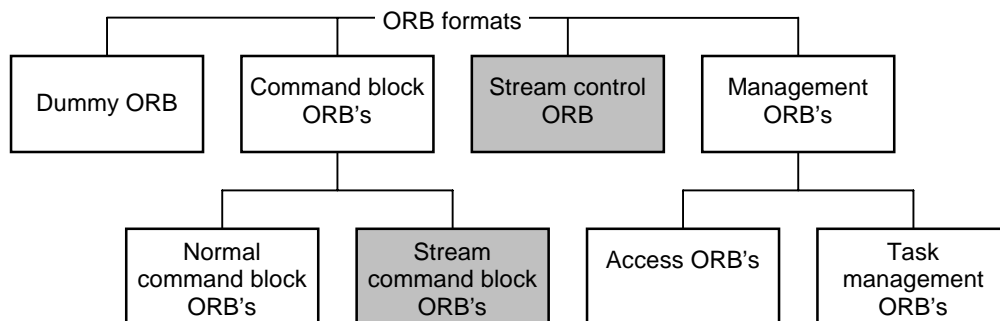


Figure 13 – ORB family tree

In the preceding figure, the ORB's that pertain to isochronous operations are shown shaded in gray. The formats of the ORB's are described in the clauses that follow. This clause specifies fields that are common to all ORB's, illustrated in the figure below.




Figure 14 – ORB format

The *notify* bit (abbreviated as *n* in the figure above) advises the target whether or not completion notification is required. When *notify* is zero, the target may elect to suppress completion notification except in the case of an error. If *notify* is one, the target shall always store a status block in initiator memory. When the target stores a status block, it shall store it at the *status_FIFO* address specified in the ORB or, if not specified in the ORB, at the address supplied in the login or create stream request. If the request completes with an error condition, the value of *notify* is ignored and a status block shall be stored at the *status_FIFO* address.

The *rq_fmt* field specifies ORB format, as defined by the table below.

Value	ORB format
0	Format specified by this standard
1	Reserved for future standardization
2	Vendor-dependent
3	Dummy (NOP) request format

The format of an ORB is uniquely determined by a combination of *rq_fmt*, the command set implemented by the target and the target agent to which the ORB is signaled. This standard specifies those parts of the ORB that are invariant across target command sets and device types.

5.1.1 Dummy ORB

Dummy ORB's may be used as placeholders within linked lists of requests. An example is the use of a dummy ORB in the initialization of a target fetch agent (see 9.1.1). The initiator shall allocate at least the maximum ORB fetch size implemented by the target. The format of a dummy ORB is illustrated below.




Figure 15 – Dummy ORB

The *next_ORB* field shall specify a null pointer or the address of an ORB and shall conform to the address pointer format illustrated by Figure 12.

The *notify* bit is as previously defined for all ORB formats.

The *rq_fmt* field is as previously defined for all ORB formats and shall be three.

An *rq_fmt* value of three is also used to indicate an ABORT TASK request to a target. See 10.4.1 for details of ORB processing by the target and for permissible completion status values.

5.1.2 Command block ORB's

Command block ORB's are used to encapsulate data transfer or device control commands for transport to the target. A target's command set and device type determine the length of the ORB's which shall be fixed for a particular command set and device type. A target reports this size in configuration ROM (see 7.4.8).

NOTE – Although device designers may select arbitrary ORB lengths, system considerations may favor some ORB sizes over others. For example, as a result of commonly implemented cache line sizes, a 32-byte ORB is appropriate for contemporary systems.

There are two kinds of command block ORB, one for normal (asynchronous) operations and one for isochronous operations. Normal command block ORB's permit the specification of a data buffer in system memory, from which or to which data is transferred by the target.

Stream command block ORB's do not specify a data buffer in system memory. Isochronous operations involve a stream of data without system memory address context. Data bytes within a stream have relative ordering with respect to each other; there is no explicit system memory address that is the source or the sink for the stream. An isochronous stream is coupled to a stream controller that can start, stop or pause the isochronous stream on Serial Bus. For this reason only a stream length governs the data transfer.

5.1.2.1 Normal command block ORB

The format of the normal command block ORB is illustrated by the figure below.

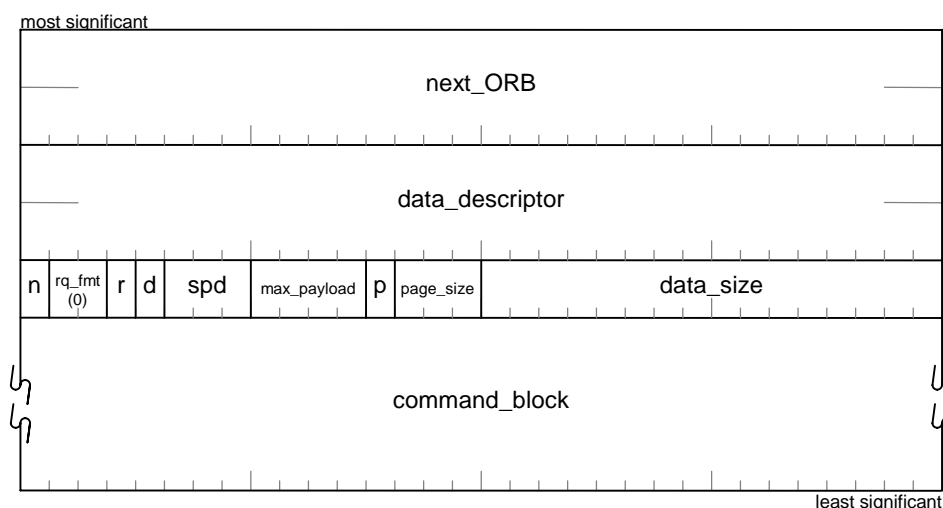


Figure 16 – Normal command block ORB

The *next_ORB* field shall specify a null pointer or the address of a dummy ORB or a normal command block ORB and shall conform to the address pointer format illustrated by Figure 12.

The *data_descriptor* field shall specify either the address of the data buffer or the address of a page table that describes the memory segments that make up the data buffer, dependant upon the value of *page_table_present* bit. The format of the *data_descriptor* field, when it directly addresses a data buffer, shall be a 64-bit Serial Bus address or, when it addresses a page table, shall be as specified by Figure 11. When *data_descriptor* specifies the address of a page table, the format of the page table shall conform to that described in 5.2.

The *notify* bit and *rq_fmt* field are as previously defined for all ORB formats. The *rq_fmt* field shall be zero.

The *direction* bit (abbreviated as *d* in the figure above) specifies direction of data transfer for the buffer. If the *direction* bit is zero, the target shall use Serial Bus read transactions to fetch data destined for the device medium. Otherwise, when the *direction* bit is one, the target shall use Serial Bus write transactions to store data obtained from the device medium.

The *spd* field specifies the speed that the target shall use for data transfer transactions addressed to the data buffer or page table, as encoded by Table 1.

Table 1 – Data transfer speeds

Value	Speed
0	S100
1	S200
2	S400
3 – 7	As standardized by P1394a

The maximum data transfer length is specified as $2^{max_payload + 2}$ bytes, which is the largest data transfer length that may be requested by the target in a single Serial Bus read or write transaction addressed to the data buffer. The *max_payload* field shall specify a maximum data transfer length less than or equal to the length permissible at the data transfer rate specified by *spd*.

The *page_table_present* bit (abbreviated as *p* in the figure above) shall be zero if *data_descriptor* directly addresses the data buffer. When *data_descriptor* indirectly addresses the data buffer, this bit shall be one.

The *page_size* field shall specify the underlying page size of the data buffer memory. A *page_size* value of zero indicates that the underlying page size is not specified. Otherwise the page size is 2^{page_size+8} bytes.

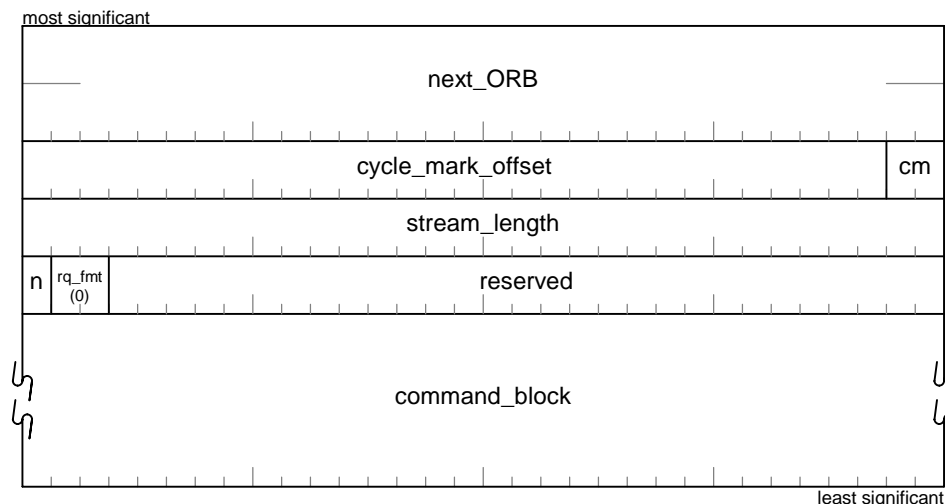
When *page_table_present* is one, the *page_size* field also specifies the format of the data structure that describes the data buffer. A *page_size* value of zero implies the unrestricted page table format (also known as a scatter/gather list). Otherwise, a nonzero *page_size* indicates a normalized page table.

If *page_table_present* is zero, the *data_size* field shall specify the size, in bytes, of the system memory addressed by the *data_descriptor* field. Otherwise *data_size* shall contain the number of elements in the page table addressed by *data_descriptor*.

The *command_block* field contains a command descriptor block not specified by this standard.

5.1.2.2 Stream command block ORB

A stream command block ORB is a structure that has the format illustrated below.

**Figure 17 – Stream command block ORB**

The *next_ORB* field shall specify a null pointer or the address of a dummy ORB or a stream command block ORB and shall conform to the address pointer format illustrated by Figure 12.

The *cm* field (together with the *cycle_mark_offset* field) specifies the location of the first quadlet of isochronous data (stream offset) as encoded by the table below.

Value	<i>cycle_mark_offset</i>	Stream offset
0	Undefined	Zero
1	Reserved (not to be used)	Reserved (not to be used)
2	Location of first CYCLE MARK	Zero
3	Location of first CYCLE MARK	<i>cycle_mark_offset</i>

The stream offset derived from the combination of *cm* and *cycle_mark_offset* specifies the location of the first quadlet of the isochronous data as an offset, in quadlets, relative to the starting medium location indicated by the *command_block*.

NOTE – The command block transported by the stream command block ORB specifies a starting location on the medium and an associated transfer length. Particularly in the case of block devices, such as mass storage, the relevant isochronous data may be a subset of the data length and may commence at a nonzero offset relative to the natural block boundaries of the medium—hence the necessity for the additional values, *stream_length* and stream offset, to completely characterize the request.

The *cycle_mark_offset* field, when *cm* has a value of two or three, specifies the location of the first CYCLE MARK packet as an offset, in quadlets, relative to the starting medium location indicated by the *command_block*.

NOTE – The *cycle_mark_offset* field may be useful to reestablish synchronization within the recorded isochronous data if a prior stream command block terminated in error.

The *stream_length* field specifies the length of data, in bytes, that is to be transferred to or from the device medium.

The *notify* bit and *rq_fmt* field are as previously defined for all ORB formats. The *rq_fmt* field shall be zero.

The *command_block* field contains a command descriptor block not specified by this standard.

5.1.3 Stream control ORB

Stream control ORB's direct the action of a logical unit stream controller. The stream controller is configured at the time of a create stream request as either a talker or a listener. When listening, the stream controller accepts isochronous data from Serial Bus in accordance with stream control ORB's, transforms the isochronous stream and then records the data on the medium as specified by isochronous requests. When talking, this process is reversed and an isochronous data stream obtained from the medium is filtered and transformed by the stream controller before isochronous packets are transmitted on Serial Bus.

The format of the stream control ORB is illustrated below.

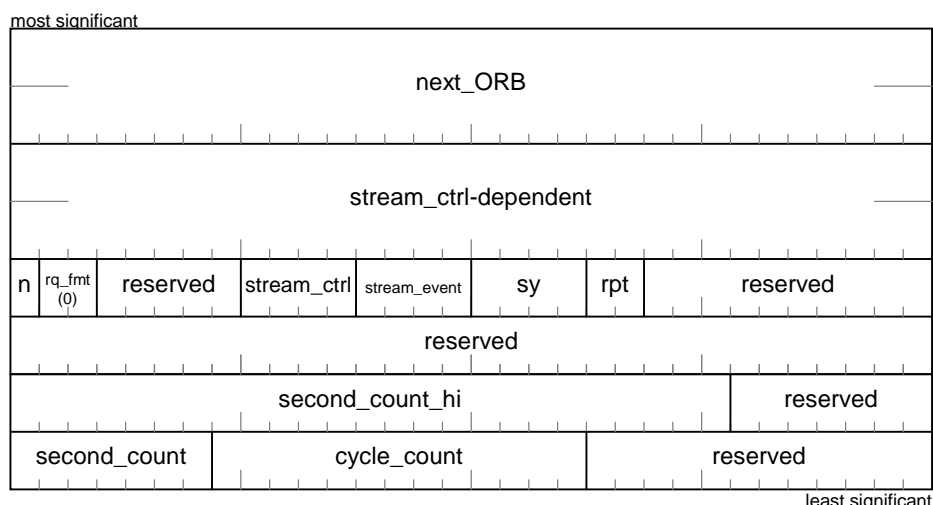


Figure 18 – Stream control ORB

The *next_ORB* field shall specify a null pointer or the address of a dummy ORB or a stream control ORB and shall conform to the address pointer format illustrated by Figure 12.

The usage of the *stream_ctrl*-dependent field varies according to the value of *stream_ctrl* and is described in more detail for each stream control function.

The *notify* bit and *rq_fmt* field are as previously defined for all ORB formats. The *rq_fmt* field shall be zero.

The *stream_ctrl* field shall specify a stream control function for the stream, as encoded below.

Value	Stream control function
0	Reserved; not to be used
1	START
2	STOP
3	PAUSE
4	UPDATE CHANNEL MASK
5	CONFIGURE CHANNELS
6	SET ERROR MODE
7	QUERY STREAM STATUS
8 – F ₁₆	Reserved for future standardization

The START control function instructs the logical unit's stream controller to commence (or resume) talking or listening on Serial Bus. The time at which the action is to occur shall be specified by the *stream_event* field in conjunction with other stream control ORB fields.

The STOP control function instructs the logical unit's stream controller to terminate the isochronous stream and to flush the stream buffers. If the target had been listening, any isochronous data already received from Serial Bus shall be made available to the isochronous commands previously enqueued at the stream command block agent. If the target had been talking, any isochronous data obtained from isochronous commands shall be discarded. The time at which the action is to occur shall be specified by the *stream_event* field in conjunction with other stream control ORB fields.

The PAUSE control function instructs the logical unit's stream controller to suspend the transfer of isochronous data with the expectation that isochronous data transfer will resume. If the target is a talker, the stream controller shall pause on the requested stream event and shall not send any isochronous packets for the stream while paused. Subject to target implementation limitations, data from isochronous commands previously enqueued at the stream command block agent may continue to accumulate at the target while the data stream is paused. If the target is a listener, the target shall pause on the requested stream event and shall discard any isochronous packets for the stream while paused. The target may flush any isochronous data already received from Serial Bus in order to make it available to any isochronous commands previously enqueued at the stream command block agent.

The UPDATE CHANNEL MASK control function instructs the logical unit's stream controller to change the set of enabled channels. The enabled channels shall be specified by the *channel_mask* field. The time at which the action is to occur shall be specified by the *stream_event* field in conjunction with other stream control ORB fields. When *stream_ctrl* specifies a value of UPDATE CHANNEL MASK, the *stream_ctrl*-dependent field shall contain a 64-bit channel mask, as shown below.




Figure 19 – Channel mask

A one in the bit position that corresponds to one of the numbered Serial Bus isochronous channels, zero to 63, indicates that the channel is to be enabled. Channel zero is represented by the most significant bit while the least significant bit represents channel 63. When a channel is enabled for listening, isochronous packets observed for that channel are transferred to the device medium under control of isochronous commands for the stream specified by *login_ID*. Conversely, when a channel is enabled for talking, an isochronous stream is obtained from the medium as directed by isochronous commands and isochronous packets are transmitted on Serial Bus for the enabled channel. The channel number specified is the channel number prior to any transformation that is a result of values in the logical unit's stream controller channel map.

The CONFIGURE CHANNELS control function instructs the logical unit's stream controller to update the 64-entry channel map maintained internally for the isochronous stream. The channel map specifies a transformation from an external channel (on Serial Bus) to an internal channel (on the medium). When listening, channel numbers observed in Serial Bus isochronous packets are replaced with numbers specified by the channel map before the isochronous data is recorded on the medium. When talking, channel numbers encountered in recorded isochronous data are replaced with numbers specified by the channel map before the isochronous packets are transmitted on Serial Bus.

NOTE – It is possible for a mapping of two or more source channels into a single destination channel to be meaningful. For example, isochronous data recorded at different times from different channels may be concatenated on the medium and subsequently replayed as a single channel.

When *stream_ctrl* specifies a value of CONFIGURE CHANNELS, the *stream_ctrl*-dependent field shall contain the address of 64-entry channel configuration map. In this case the *stream_ctrl*-dependent field shall conform to the format for address pointers specified by Figure 11 and shall address the same node as the initiator; consequently the *node_ID* field of this address pointer shall be reserved.

The channel configuration map consists of 64 quadlet entries. The channel configuration map is indexed by the source channel. When the target is a listener, the source channel is that observed in the isochronous packet header on Serial Bus. When the target is a talker, the source channel is that

previously recorded on the medium. The format of the channel configuration map entries is illustrated below.




Figure 20 – Channel configuration map entry

The *spd* field is valid only if the target is configured as a talker. In this case *spd* determines the speed at which the isochronous packets shall be transmitted, as encoded by Table 1.

The *dest_channel* field shall specify the channel number transformation for either listening or talking. When the target is a listener, the observed channel number in the isochronous packet header shall be replaced with *dest_channel* as the data is recorded on the medium. When the target is a talker, the channel number obtained from the medium shall be replaced with *dest_channel* at the time the isochronous packet is transmitted.

The *sid* field is valid only if the target is configured as a talker. If the recorded isochronous data conforms to the CIP format specified in 11.5, the *sid* field in the previously recorded CIP header shall be replaced with the *sid* value specified by the channel map when the packet is transmitted.

The SET ERROR MODE control function instructs the logical unit's stream controller to configure its error handling mode as specified by the *rpt* field.

The QUERY STREAM STATUS control function instructs the logical unit's stream controller to return status information that indicates whether or not the stream controller is ready to accept a START control function.

NOTE – Assume that a target is to be instructed to listen to isochronous data and transfer the stream to device medium. If the starting medium location is at a nonzero offset relative to a block boundary, some implementations may require time to read previously recorded data from the medium before being ready to commence recording the new isochronous data. Subsequent to enqueueing an isochronous command at the stream command block agent, the QUERY STREAM STATUS control function may be used to determine if the target is ready to accept a START control function.

The *stream_event* field is valid only if the *stream_ctrl* field specifies a value of START, STOP, PAUSE or UPDATE CHANNEL MASK. When one of these control functions is specified, the *stream_event* field specifies the time at which the action is to take place, as encoded below.

Value	Stream event code
0	IMMEDIATE
1	CYCLE MATCH
2	SY MATCH
3	FIRST DATA
4 – F ₁₆	Reserved for future standardization

A value of IMMEDIATE instructs the logical unit's stream controller to perform the specified action as soon as possible, within the capabilities of the target implementation.

A value of CYCLE MATCH instructs the logical unit's stream controller to perform the specified action at the cycle time specified by *second_count_hi*, *second_count* and *cycle_count*.

A value of SY MATCH instructs the logical unit's stream controller to perform the specified action on the isochronous cycle for which the *sy* field of an isochronous packet for any enabled channel matches the *sy* field in the stream control ORB. A *stream_event* value of SY MATCH is valid only if the logical unit's stream controller is configured as a listener.

A value of FIRST DATA instructs the logical unit's stream controller to perform the specified action when isochronous data is observed for any enabled isochronous channel. A *stream_event* value of FIRST DATA is valid only if the logical unit's stream controller is configured as a listener.

NOTE – A *stream_event* field value of FIRST DATA may have effects similar to IMMEDIATE, in that it is possible for isochronous data to be recorded immediately. The difference between the two stream events is apparent if no isochronous packets for any of the enabled channels are present when the stream control ORB is executed. If IMMEDIATE is specified, CYCLE MARK packets are recorded as each cycle start is observed. If FIRST DATA is specified, no packets are recorded until the first isochronous packet for an enabled channel is observed. When this event occurs, a CYCLE MARK packet with the most recent cycle start data is recorded followed by a DATA packet for the enabled channel.

The *sy* field is valid only if the logical unit's stream controller is configured as a listener and the *stream_event* field specifies SY MATCH. See the preceding description of *stream_event*.

The *rpt* field specifies an operational mode for the logical unit's stream controller, as described in the table below. The *rpt* field is valid only if the *stream_ctrl* field specifies SET ERROR MODE.

Value	Error handling mode
0	Report errors and halt stream
1	Report errors and continue stream
2	Ignore all errors
3	Reserved for future standardization

Different sorts of errors may be detected when the logical unit's stream controller is configured as a talker or a listener. If an error occurs, the stream controller shall take one of three actions, as specified by the value of *rpt*:

- Report the error by writing unsolicited status to the initiator and then halting isochronous data transfers by performing the equivalent of a STOP control function with a *stream_event* value of IMMEDIATE;
- Report the error by writing unsolicited status to the initiator but continue isochronous data transfers;
or
- Ignore the error and continue isochronous data transfers.

A more detailed description of isochronous errors and how they are handled is provided in 12.3.

The *second_count_hi*, *second_count* and *cycle_count* fields are valid only if the *stream_ctrl* field specifies START, STOP, PAUSE or UPDATE CHANNEL MASK and the *stream_event* field specifies CYCLE MATCH. Together, these fields specify a cycle time for comparison with the target's cycle clock. An equal comparison occurs if the *second_count_hi* field matches the field of the same name in the target's BUS_TIME register and if both the *second_count* and *cycle_count* fields match their corresponding fields in the target's CYCLE_TIME register.

5.1.4 Management ORB's

Management ORB's are 32-byte data structures that encapsulate several types of management request:

- access requests (which include login and logout requests); and
- task management requests.

Unlike the normal command block, stream command block and stream control ORB's (which are all implicitly associated with a particular task set or stream by virtue of the fetch agent to which they are addressed), most management ORB's explicitly declare the task set or stream for which they are intended.

Management ORB's have the general format illustrated below. Note that since they lack a *next_ORB* field, they cannot be linked together to form a list.

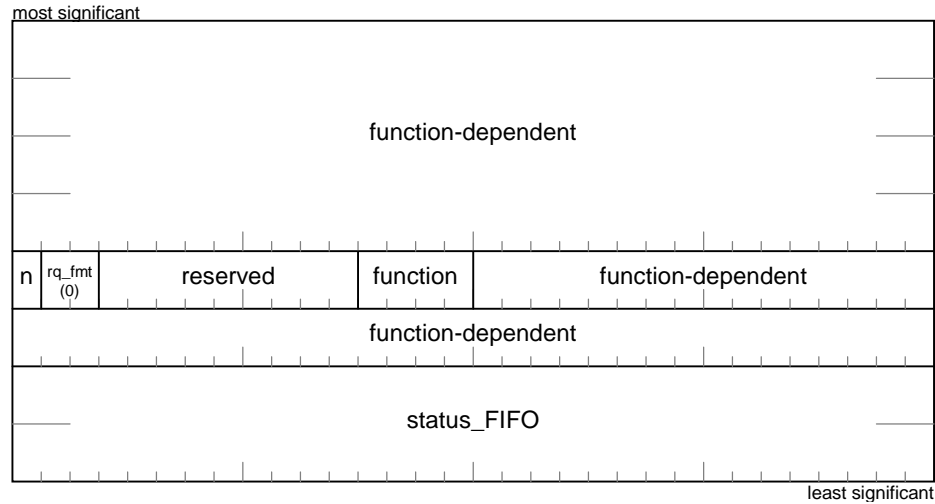


Figure 21 – Management ORB

The *notify* bit and *rq_fmt* field are as previously defined for all ORB formats. The *rq_fmt* field shall be zero.

The *function* field specifies the management function requested, as defined by the table below.

Table 2 – Management request functions

Value	Management function
0	LOGIN
1	QUERY LOGINS
2	CREATE STREAM
3	RECONNECT
4	SET PASSWORD (see Annex C)
5 – 6	Reserved for future standardization
7	LOGOUT
8 – 9	Reserved for future standardization
A ₁₆	TERMINATE TASK
B ₁₆	ABORT TASK
C ₁₆	ABORT TASK SET
D ₁₆	CLEAR TASK SET
E ₁₆	LOGICAL UNIT RESET
F ₁₆	TARGET RESET

The *status_FIFO* field shall specify an address allocated for the return of status information generated by the management request. The *status_FIFO* field shall conform to the format for address pointers specified by Figure 11 and shall address the same node as the initiator; consequently the *node_ID* field of this address pointer shall be reserved.

NOTE – The *status_FIFO* address explicitly specified within a management ORB may differ from the status FIFO address implicitly associated with command block requests (normal or stream) or stream control requests. The address in those cases is established by a LOGIN or CREATE STREAM request and is not altered by other management requests.

5.1.4.1 Login ORB

Before any other requests (except QUERY LOGINS) can be made of a target, the initiator shall first complete a login procedure that uses the ORB format shown below.

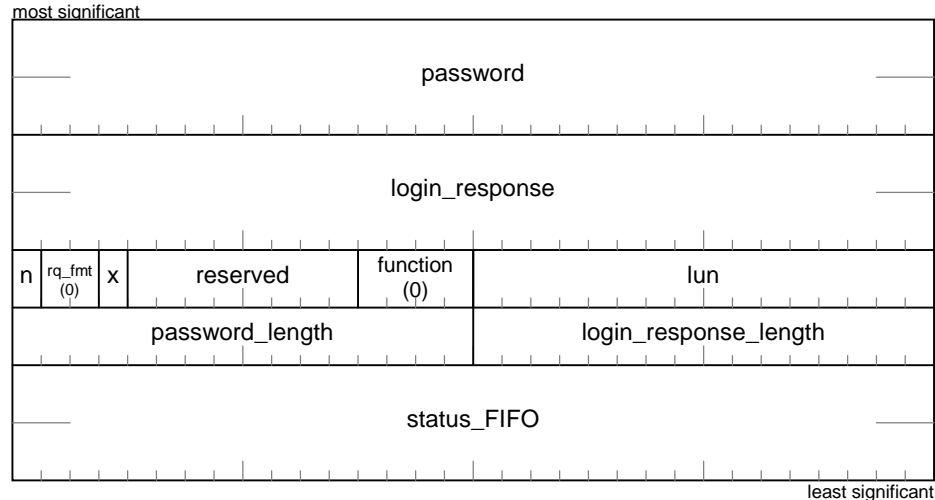


Figure 22 – Login ORB

The *password* and *password_length* fields may specify optional, command set-dependent information used to validate the login request. If *password_length* is zero, the *password* field may contain immediate data. When *password_length* is nonzero, the *password* field shall conform to the format for address pointers specified by Figure 11 and shall specify the address of a buffer. The buffer shall be in the same node as the initiator and shall be accessible to a Serial Bus block read request with a data transfer length less than or equal to *password_length*. The format and usage of password data, whether immediate or indirectly addressed, are beyond the scope of this standard.

The *login_response* and *login_response_length* fields specify the address and size of a buffer allocated for the return of the login response. The *login_response* field shall conform to the format for address pointers specified by Figure 11. The buffer shall be in the same node as the initiator and shall be accessible to a Serial Bus block write request with a data transfer length less than or equal to *login_response_length*. If the status block stored at the *status_FIFO* address indicates an unsuccessful login, no response data shall be stored.

The *notify* bit and the *rq_fmt* field are as previously defined for management ORB formats.

The *exclusive* bit (abbreviated as *x* in the figure above) shall specify target behavior with respect to concurrent login to a logical unit. When *exclusive* is zero, the target, subject to its own implementation capabilities, may permit more than one initiator to login to a logical unit. If *exclusive* is one the target shall permit only one login to a logical unit at a time; see 8.2.1 for a description of target behavior.

The *lun* field specifies the logical unit number (LUN) to which the request is addressed.

The *status_FIFO* field shall specify an address allocated for the return of status for the LOGIN request, status for all subsequent requests signaled to the *command_block_agent* allocated for this login and any unsolicited status generated by the logical unit.

Upon successful completion of a login, the response is returned in the format illustrated below.




Figure 23 – Login response

The *length* field shall specify the length, in bytes, of the login response data. If *login_response_length* in the login request is too small for the transfer of all the login response data, the *length* field shall not be adjusted to reflect the truncation.

The initiator shall use the *login_ID* value returned by the target to identify all subsequent requests directed to the target's management agent that pertain to this login.

The *command_block_agent* field specifies the base address of the agent's CSR's, which are defined in 6.4. This field shall conform to the format for address pointers specified by Figure 11. The *node_ID* portion of the field shall have a value equal to the most significant 16 bits of the target's NODE_IDS register.

5.1.4.2 Query logins ORB

An initiator may determine the EUI-64 and node ID of all currently logged-in initiators by means of a query logins request, whose format is illustrated below.

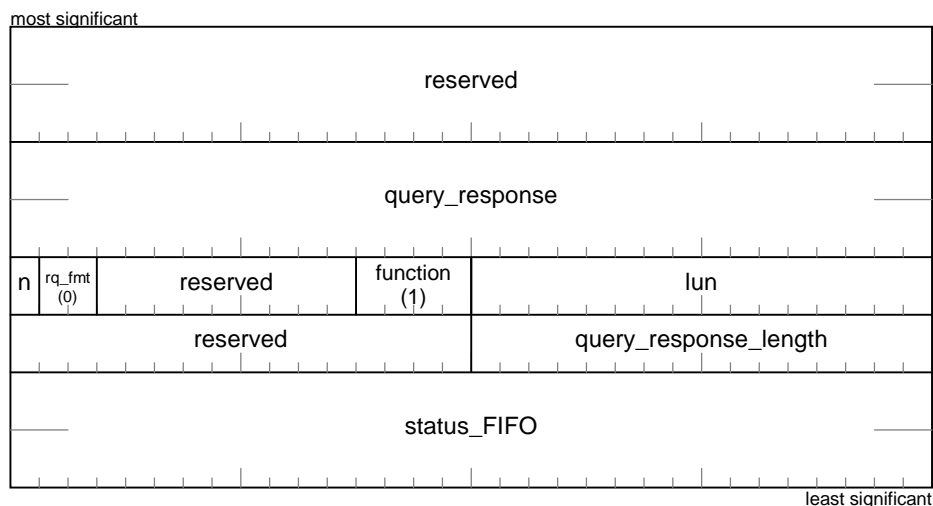


Figure 24 – Query logins ORB

The *query_response* and *query_response_length* fields specify the address and size of a buffer for the return of the query results. The buffer shall be accessible to a Serial Bus block write request with a data transfer length less than or equal to *query_response_length*.

The *notify* bit, *rq_fmt* and *status_FIFO* fields are as previously defined for management ORB formats.

The *query_response* and *status_FIFO* fields shall conform to the format for address pointers specified by Figure 11. All of these buffers shall be in the same node as the initiator; consequently the *node_ID* field of these addresses shall be reserved.

The *lun* field specifies the logical unit number (LUN) to which the request is addressed.

The query response data returned shall have the following format.

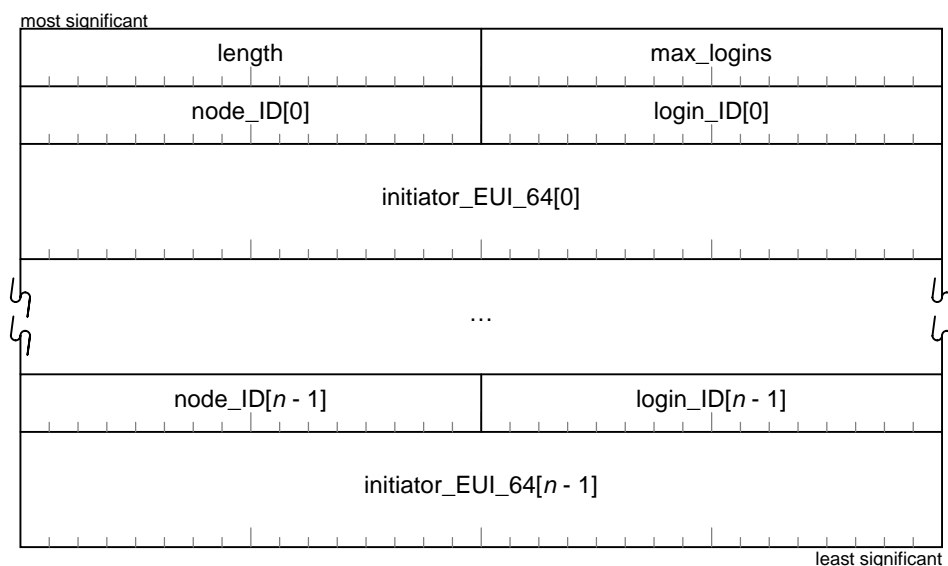


Figure 25 – Query logins response format

The *length* field shall specify the length, in bytes, of the query response data. If *query_response_length* in the query logins request is too small for the transfer of all the query response data, the *length* field shall not be adjusted to reflect the truncation. The value of the *length* field shall be equal to $4 + 12 * n$, where n is the number of logged-in initiators.

The *max_logins* field shall specify the maximum concurrent logins that may be accepted by the logical unit.

The remainder of the query response is a variable-length array of 12-byte entries, one for each logged-in initiator, each of which contains a *node_ID*, *login_ID* and *initiator_EUI_64* field.

The *node_ID* field of an entry shall specify the node ID of a logged-in initiator. If a Serial bus reset has occurred since the login was established and the initiator has not reconnected the login, the *node_ID* field shall have a value of $FFFF_{16}$.

The *login_ID* field of an entry shall specify the login ID provided to the initiator as a result of its successful login.

The *initiator_EUI_64* field of an entry shall specify the EUI-64 obtained by the target from the initiator's configuration ROM at the time the login was validated.

5.1.4.3 Create stream ORB

Before any stream requests can be made of a target, the initiator shall first complete a create stream procedure that uses the ORB format shown below.

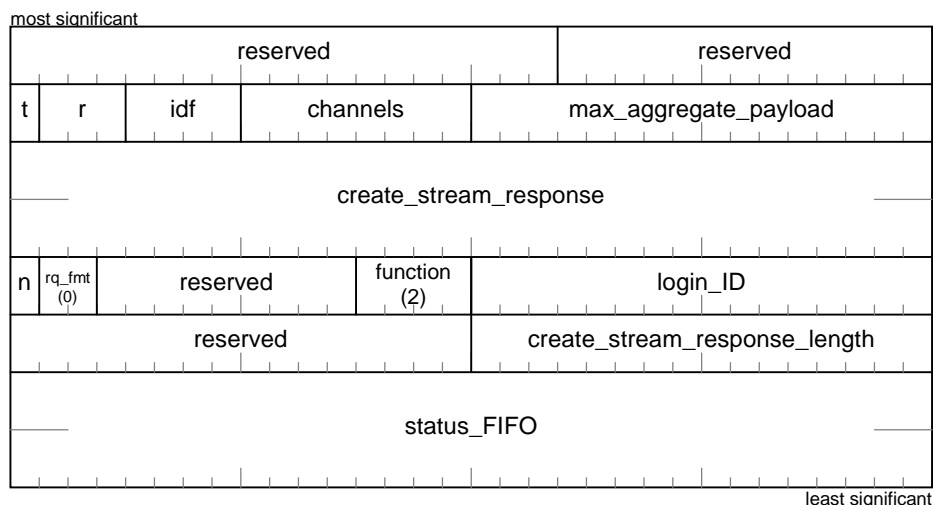


Figure 26 – Create stream ORB

The *talker* bit (abbreviated as *t* in the figure above) shall specify the type of isochronous stream requested. If the target resources are to be configured for listening, *talker* shall be zero.

The *idf* field specifies the manner in which isochronous data is recorded or played back. Valid values for *idf* are encoded in the table below.

Value	Name	Description	Reference
0	Unstructured	The data is an unstructured stream; cycle marks are not present in the recorded data format	
1	—	Reserved for future standardization	
2	Structured (without cycle mark indices)	The data is structured by cycle marks, which are recorded on the medium.	Section 11 (also 12.2.3)
3	Structured (with cycle mark indices)	The data is structured by cycle marks which are recorded on the medium; a cycle mark index is present every 512 bytes.	Section 11 (also 11.4 and 12.2.3)
4 – F ₁₆	—	Reserved for future standardization	

The *channels* field specifies the maximum number of isochronous channels that are to be simultaneously transmitted or received.

The *max_aggregate_payload* field is the aggregate maximum isochronous payload that the target is requested to support for the stream. That is, the sum of all the *data_length* fields of Serial Bus isochronous packets transmitted or received for all of the stream's isochronous channels shall not exceed *max_aggregate_payload* in a single isochronous cycle.

The *create_stream_response* and *create_stream_response_length* fields specify the address and size of a buffer allocated for the return of the create stream response. The *create_stream_response* field shall conform to the format for address pointers specified by Figure 11. The buffer shall be in the same node as the initiator and shall be accessible to a Serial Bus block write transaction with a data transfer length less

than or equal to *create_stream_response_length*. If the status block stored at the *status_FIFO* address indicates an unsuccessful attempt to create a stream, no response data shall have been stored.

The *notify* bit and the *rq_fmt* field are as previously defined for management ORB formats.

The *login_ID* field shall specify a login ID value obtained as the result of a successful login.

The *status_FIFO* field shall specify an address allocated for the return of status for the CREATE STREAM request, status for all subsequent requests signaled to either the *command_block_agent* or *stream_control_agent* allocated for this login and any unsolicited status generated by the logical unit for this stream.

Upon successful completion of a create stream request, the response is returned in the format illustrated below.

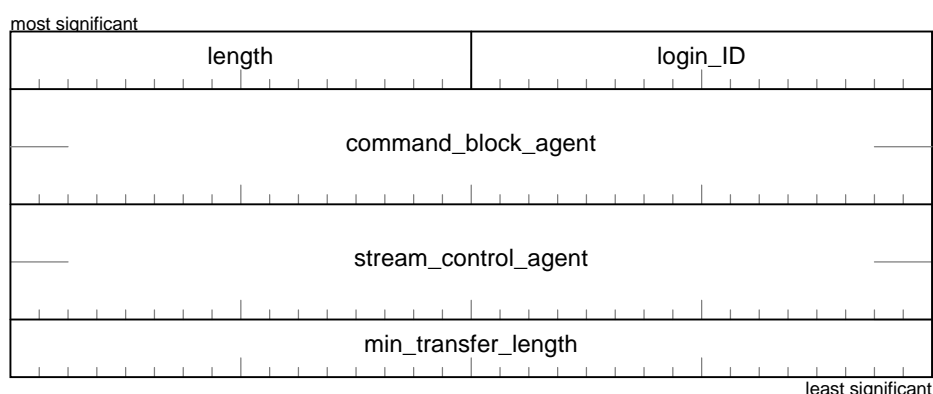


Figure 27 – Create stream response

The *login_ID* identifies an isochronous stream for which target resources have been allocated. The initiator shall use this value to identify all subsequent requests directed to the target's management agent that pertain to this stream.

The *command_block_agent* and *stream_control_agent* fields specify the base address of the agent's CSR's, which are defined in 6.4. Both fields shall conform to the format for address pointers specified by Figure 11. The *node_ID* portion of either field shall have a value equal to the most significant 16 bits of the target's *NODE_IDS* register.

The *min_transfer_length* field advises the initiator of the minimum *stream_length* value desired by the target in stream command block ORB's in order to sustain the isochronous data transfer rate requested by the login. If the initiator presents any stream command block ORB's whose *stream_length* value is less than this minimum, the target may experience underflow or overflow in isochronous data while talking or listening at the requested rate. Even if this minimum is respected it is still possible for underflow or overflow to occur.

5.1.4.4 Reconnect ORB

After a Serial Bus reset it is possible for an initiator's 16-bit node ID to change. Since a target validates all writes to agent CSR's by the node ID of the initiator, an initiator shall reestablish validated access before it may signal new requests to the target. This is accomplished by means of a reconnect request, with the format shown below.

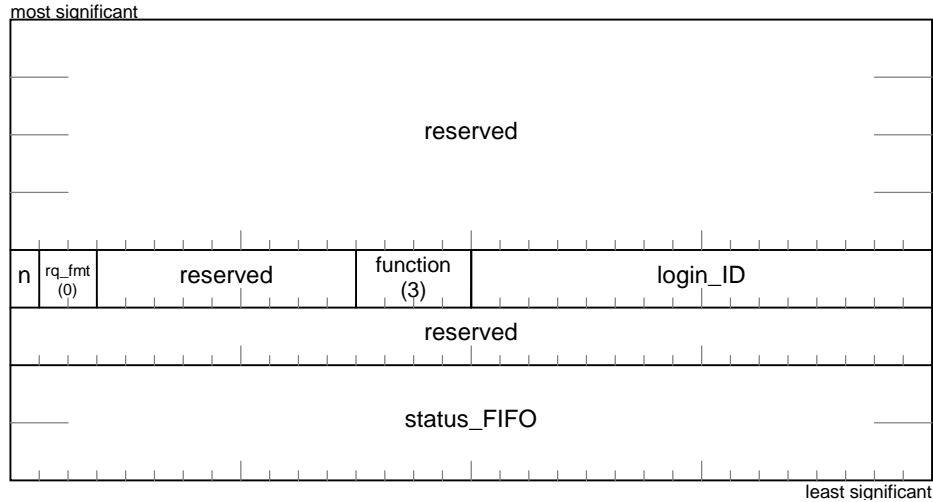


Figure 28 – Reconnect ORB

The *notify* bit and the *rq_fmt* field are as previously defined for management ORB formats.

The *login_ID* field shall specify a login ID value obtained as the result of a successful login. The target shall verify that the EUI-64 of the initiator requesting the login reestablishment matches the EUI-64 previously saved by the target for the *login_ID*.

The *status_FIFO* field shall specify an address allocated for the return of status for the RECONNECT request, only. The contents of this field shall not update the status FIFO address established by the successful login that returned *login_ID*.

Upon successful reestablishment of the login, the initiator may signal requests to the target agent at the same CSR addresses returned in the original login response data. The initiator shall also use the *login_ID* value to identify all requests directed to the target's management agent that pertain to the reestablished login.

Any isochronous streams established with the same *login_ID* value specified in the reconnect ORB are also reestablished. The login ID's of the isochronous streams remain the same.

5.1.4.5 Logout ORB

When an initiator wishes to relinquish its access privileges for a logical unit or an isochronous stream, it shall perform a logout with the ORB format shown below.

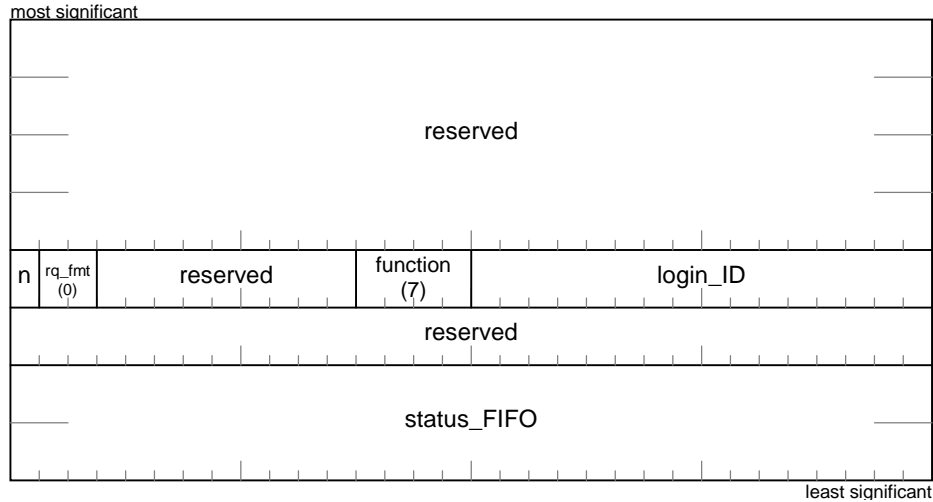


Figure 29 – Logout ORB

The *notify* bit, *rq_fmt* and *status_FIFO* fields are as previously defined for management ORB formats.

The *login_ID* field shall specify a login ID value obtained as the result of a successful login or create stream request.

5.1.4.6 Task management ORB

The task management ORB is used to control task sets. This ORB shall have the format defined below.

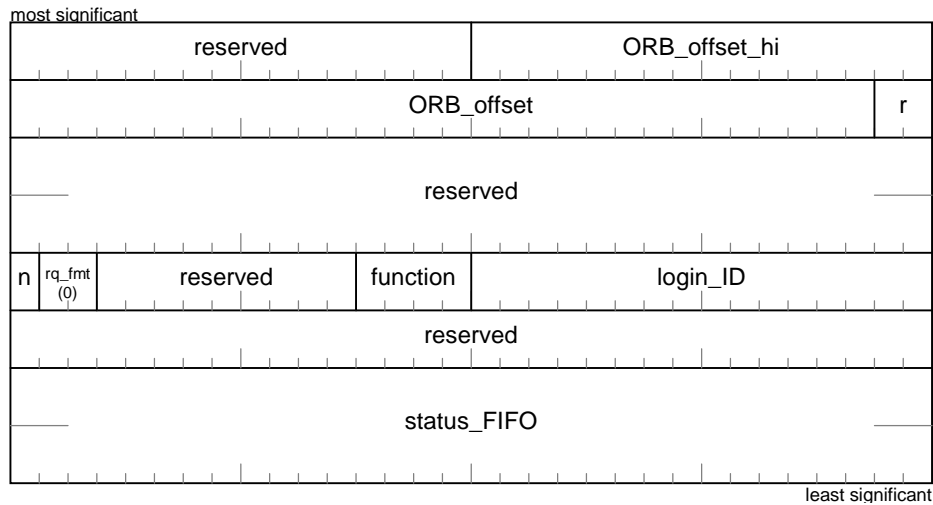


Figure 30 – Task management ORB

The *ORB_offset_hi* and *ORB_offset_lo* fields together form the *ORB_offset* field, which identifies the task to which the management function applies. *ORB_offset* is derived by taking the least significant 48 bits of the Serial Bus address of the ORB and discarding the least significant two bits. The *ORB_offset* field is ignored unless the *function* field is TERMINATE TASK or ABORT TASK. All tasks are uniquely identified by their Serial Bus address of the ORB that initiated the task.

The *notify* bit, *rq_fmt* and *status_FIFO* fields are as previously defined for management ORB formats.

The *function* field shall contain a value of TERMINATE TASK, ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET or TARGET RESET, as defined by Table 2.

The *login_ID* shall be set to the value returned in the login response data for the task set to which the task management request is directed.

NOTE – In the case of TARGET RESET, which does not pertain to any one task set, *login_ID* shall be set to a value obtained as the result of any successful login completed by the initiator.

5.2 Page tables

The data buffer specified by a normal command block ORB is described by the *data_descriptor*, *page_table_present*, *page_size* and *data_size* fields. The data buffer is a logically contiguous area in system memory. As previously described, when *page_table_present* is zero, the data buffer is also contiguous within Serial Bus address space and no more than 65,535 bytes in length. In this case, *data_descriptor* contains the 64-bit address of the data buffer and *data_size* specifies its length, in bytes.

When the data buffer cannot be directly addressed (either because it is discontinuous or too large), it is necessary to describe it *via* a page table. A page table is a variable-length array of elements, each of which describes a segment that is contiguous within Serial Bus address space. Page table elements shall be octlet aligned.

The presence of a page table is indicated by the value of *page_table_present* in the ORB. When *page_table_present* is one, the *data_descriptor* field in the ORB shall contain the address of the page table and the *data_size* field shall contain the number of elements in the page table.

Page tables may have one of two formats: an unrestricted page table or a normalized page table. The page table format is determined by *page_size*. When *page_size* is zero there are no underlying page boundaries to restrict the size or alignment of data buffer segments; this is the unrestricted format. Otherwise the size and alignment of data buffer segments is determined by the nonzero *page_size*; this is the normalized format.

When a page table is used it shall be located in the same node as the data buffer it describes. The *spd* and *max_payload* fields of the ORB shall describe data transfer capabilities for both the data buffer and the page table. System memory addressed by a target request subaction that accesses the data buffer shall be entirely contained within a data buffer segment described by a single page table element.

5.2.1 Unrestricted page tables

An unrestricted page table shall be contiguous within Serial Bus address space and shall be accessible to block read requests with a *data_length* less than or equal to $data_size * 8$ bytes. The format of elements in an unrestricted page table is shown by Figure 31.




Figure 31 – Page table element (unrestricted page table)

The *segment_length* field shall specify the length, in bytes, of the portion of the data buffer (segment) described by the page table element. The value of *segment_length* shall be nonzero.

The *segment_base_hi* and *segment_base_lo* fields together shall specify the base address of the segment within the node's 48-bit system memory address range.

The 64-bit system memory address used to address the data is formed by the concatenation of the 16-bit *node_ID* field from the *data_descriptor* field in the ORB, *segment_base_hi* and *segment_base_lo*.

5.2.2 Normalized page tables

A normalized page table shall be contiguous within Serial Bus address space and shall be accessible to Serial Bus block read transactions with a *data_length* less than or equal to 2^{page_size+8} bytes so long as they do not cross Serial Bus address boundaries that occur every 2^{page_size+8} bytes.




Figure 32 – Page table element (when *page_size* equals four)

NOTE – In the figure above, the field widths of *segment_base_lo* and *segment_offset*, 20 and 12 bits, respectively, are chosen only for the purposes of illustration. The size of *segment_base_lo* and *segment_offset* vary according to *page_size*. The field width, in bits, of *segment_offset* shall be *page_size*+8. In the example shown above, the page size is assumed to be 4096 bytes.

The *segment_length* field shall specify the length, in bytes, of the portion of the data buffer (segment) described by the page table element. The value of *segment_length* shall be less than or equal to 2^{page_size+8} .

The *segment_base_hi* and *segment_base_lo* fields together shall specify the base address of the segment within the node's 48-bit system memory address range.

The *segment_offset* field shall specify the starting address for data transfer within the segment.

The 64-bit system memory address used to address the data is formed by the concatenation of the 16-bit *node_ID* field from the *data_descriptor* field in the ORB, *segment_base_hi*, *segment_base_lo* and *segment_offset*.

In all page table elements, the sum of *segment_length* and *segment_offset* shall be less than or equal to 2^{page_size+8} .

In addition to the preceding requirements, the values of *segment_length* and *segment_offset* are constrained by their position within the page table. These additional restrictions are summarized below.

Element	Total page table elements		
	1	2	n (where $n \geq 3$)
0	No additional restrictions	$segment_length = 2^{page_size+8} - segment_offset$	
$1 - n-2$	—	$segment_offset = 0$	$segment_offset = 0$ $segment_length = 2^{page_size+8}$
$n-1$	—	—	$segment_offset = 0$

5.3 Status block

Upon completion of a request, if the *notify* bit in the ORB is one or if there is error status to report, the target shall store all or part of the status block shown below. If the *status_FIFO* address is explicitly provided as part of the ORB to which the status pertains, the target shall store the status block at the address specified. Otherwise, the target shall store the status block at the *status_FIFO* address provided by the initiator as part of the login or create stream request. The target may also store unsolicited status at this address, as defined in 9.4.

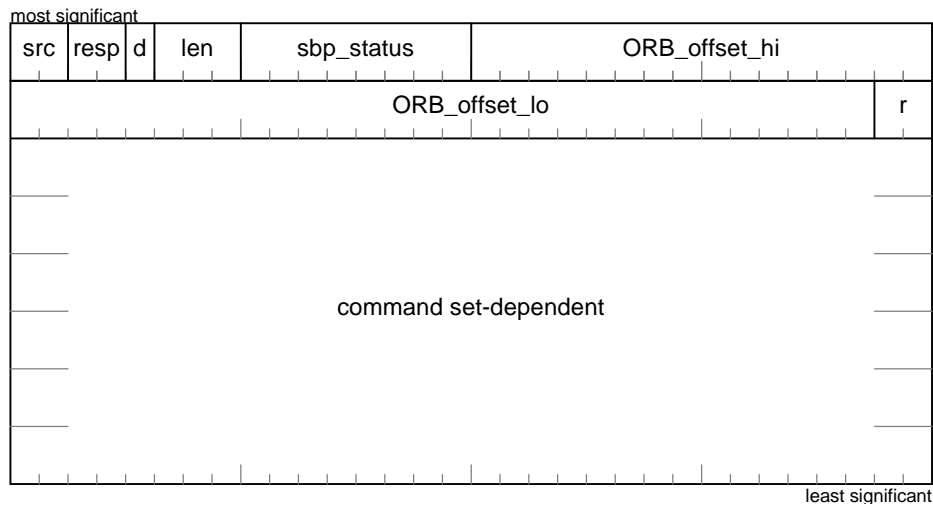


Figure 33 – Status block format

The target shall store a minimum of eight bytes of status information and may store up to the entire 32 bytes defined above so long as the amount of data stored is an integral number of quadlets. A truncated status block shall be interpreted as if the omitted fields had been stored as zeros. The target shall use a single Serial Bus block write transaction to store the status block at the *status_FIFO* address.

The *src* field indicates the origin of the status block, as encoded by the table below.

Value	Description
0	The status block pertains to an ORB identified by <i>ORB_offset</i> ; at the time the ORB was most recently fetched by the target the <i>next_ORB</i> field did not contain a null pointer.
1	The status block pertains to an ORB identified by <i>ORB_offset</i> ; at the time the ORB was most recently fetched by the target the <i>next_ORB</i> field was null.
2	The status block is unsolicited and contains device status information; the contents of the <i>ORB_offset</i> field shall be ignored.
3	The status block is unsolicited and contains isochronous error report information as specified by 12.3.

The *resp* field shall specify the SBP-2 response status for the request identified by *ORB_offset*. Response values are encoded by *resp* as shown by the table below.

Value	Name	Description
0	REQUEST COMPLETE	The request completed without transport protocol error (Either <i>sbp_status</i> or command set-dependent status information may indicate the success or failure of the request)
1	TRANSPORT FAILURE	The target detected a nonrecoverable transport failure that prevented the completion of the request
2	ILLEGAL REQUEST	There is an unsupported field or bit value in the ORB; the <i>sbp_status</i> field may provide additional information
3	VENDOR DEPENDENT	The meaning of <i>sbp_status</i> shall be specified by the vendor

The *dead* bit (abbreviated as *d* in the figure above) shall indicate whether or not the target fetch agent transitioned to the dead state upon storing the status block. When *dead* is zero, the reported status has not affected the state of the fetch agent. If the *dead* bit is set to one, the fetch agent transitioned to the dead state as a consequence of the error condition reported by the status block.

The *len* field shall specify the quantity of valid status block information stored at the *status_FIFO* address. The size of the status block is encoded as *len* + 1 quadlets.

The *sbp_status* field provides additional information that qualifies the response status in *resp*. The meanings assigned to *sbp_status* vary according to the value of *resp* and are described below.

When *resp* is equal to zero, REQUEST COMPLETE, the possible values for *sbp_status* are specified by the table below. Any value not enumerated is reserved for future standardization.

Value	Description
0	No additional sense to report
1	Request type not supported
2	Speed not supported
3	Page size not supported
4	Access denied
5	Logical unit not supported
6	Maximum payload too small
7	Too many channels
8	Resources unavailable
9	Function rejected
10	Login ID not recognized
11	Dummy ORB completed
12	Request aborted
FF ₁₆	Unspecified error

If a Serial Bus error occurs in the transport (*resp* is equal to one, TRANSPORT FAILURE), the *sbp_status* field either shall have a value of FF₁₆, unspecified error, or else the field shall be redefined as illustrated below. This format provides for the return of additional information about the transport failure.




Figure 29 – TRANSPORT FAILURE format for *sbp_status*

The *object* field shall specify which component of an SBP-2 request, the ORB, the data buffer or the page table, was referenced by the target when the error occurred. The value of *object* shall be as defined by the following table.

Value	Referenced object
0	Operation request block (ORB)
1	Data buffer
2	Page table
3	Unable to specify

The *serial_bus_error* field shall specify the error response for the failed request, as encoded by the table below.

Value	Serial Bus error	Comment
0	Missing acknowledge	
1	Reserved; not to be used	
2	Time-out error	An <i>ack_pending</i> was received for the request but no response was received within the time-out limit
3	Reserved; not to be used	
4 – 6	Busy retry limit exceeded	The value reflects the last acknowledge, <i>ack_busy_X</i> , <i>ack_busy_A</i> or <i>ack_busy_B</i> .
7 – B ₁₆	Reserved for future standardization	
C ₁₆	Conflict error	A resource conflict was detected by the addressed node
D ₁₆	Data error	The data field failed the CRC check or the observed length of the payload did not match the <i>data_length</i> field
E ₁₆	Type error	A field in the request was set to an unsupported value or an invalid transaction was attempted (e.g., a write to a read-only address)
F ₁₆	Address error	The <i>destination_offset</i> field specified an inaccessible address in the addressed node

In the cases of conflict error and data error, these are errors that the target may retry up to an implementation-dependent limit before reporting TRANSPORT FAILURE.

No additional information is provided in *sbp_status* when *resp* equals two, ILLEGAL REQUEST. In this case, *sbp_status* shall be set to FF₁₆.

NOTE – An SBP-2 response code of ILLEGAL REQUEST shall not be used to indicate unsupported fields or bit values in the command set-dependent portion of the ORB. This response code shall be used only to indicate an error in the first 20 bytes of the ORB.

The *ORB_offset_hi* and *ORB_offset_lo* fields together form an *ORB_offset* field that uniquely identifies the ORB to which the status block pertains. If *src* is zero or one, the target shall form *ORB_offset* from the least significant 48 bits of the Serial Bus address used to fetch the ORB; the least significant two bits shall be discarded. When the status block contains unsolicited device status (*src* is equal to two), the *ORB_offset* field shall be ignored by the initiator. Otherwise, when *src* equals three, the status block contains an isochronous error report in the format illustrated below.




Figure 30 – Unsolicited status format for isochronous errors

The *src* field shall be three to indicate unsolicited status that describes an isochronous error.

The *resp* field shall specify a response status of TRANSPORT FAILURE.

The *len* field shall be equal to one.

The *stream_error* field shall specify the nature of the isochronous error, as encoded by the table below.

Value	Stream error description
0	Reserved (not to be used)
1	Missing CYCLE START packet
2	Data CRC error in received isochronous packet
3	Data length error in received isochronous packet
4	Internal underflow with the result that recorded isochronous data was not transmitted on Serial Bus
5	Internal overflow with the result that isochronous data observed on Serial Bus was not recorded on the medium
6 — FE ₁₆	Reserved for future standardization
FF ₁₆	Unspecified error

The *seconds* field shall contain the least significant 19 bits of the BUS_TIME register at the time of the isochronous stream error.

The *cycle_count* field shall contain the cycle count, between zero and 7999, at the time of the error. The cycle count shall be obtained from the target's free-running cycle timer and shall not be latched from the last observed CYCLE START packet.

For all status block formats, the remainder of the status block after the first two quadlets, up to an overall maximum of 32 bytes, is command set-dependent.

6 Control and status registers

The control and status registers (CSR's) implemented by a target shall conform to the requirements defined by this standard and its normative references. The CSR's are arranged in four principal categories:

- core registers required by ISO/IEC 13213:1994;
- bus-dependent registers required by IEEE Std 1394-1995; and
- unit architecture registers required by this standard.

Unless otherwise specified, all registers shall support quadlet read and quadlet write transactions. The registers defined in 6.3 and 6.4 shall ignore broadcast write requests.

6.1 Core registers

The CSR architecture standardizes the locations and functions of core registers. The addresses of these registers are specified in terms of offsets, in bytes, within initial register space, where the base address of initial register space is FFFF F000 0000₁₆ relative to initial node space. IEEE Std 1394-1995 should be consulted for detailed descriptions of these core registers; the table below summarizes which core registers are mandatory for targets.

Offset	Register name	Description
0	STATE_CLEAR	State and control information
4	STATE_SET	Sets STATE_CLEAR bits
8	NODE_IDS	Contains the 16-bit <i>node_ID</i> value used to address the node
0C ₁₆	RESET_START	Resets the node's state
18 ₁₆ – 1C ₁₆	SPLIT_TIMEOUT	Time limit for split transactions

The CSR architecture and IEEE Std 1394-1995 broadly define the effects of a write to the RESET_START register. In addition to those requirements, a write to RESET_START should cause all of a node's SBP-2 units to reset in the same fashion as a power reset.

NOTE – Because of the potential for malicious interference in target operations by an unauthorized node, it is recommended that a write to RESET_START have no effect upon a target unless either a) there are no logged-in initiators or b) the *source_ID* of the write matches that of one of the currently logged-in initiators.

6.2 Serial Bus-dependent registers

The CSR architecture reserves a portion of initial register space for bus-dependent uses. Serial Bus defines registers within this address space, whose addresses are specified in terms of offsets, in bytes, within initial register space, where the base address of initial register space is FFFF F000 0000₁₆ relative to initial node space. IEEE Std 1394-1995 should be consulted for detailed descriptions of these core registers; the table below summarizes which Serial Bus-dependent registers are mandatory for targets.

Offset	Register name	Description
210 ₁₆	BUSY_TIMEOUT	Controls transaction layer retry protocols

Isochronous capabilities are optional for targets. If a target supports isochronous operations, it shall be cycle master capable and isochronous resource manager capable as well as isochronous capable. These capabilities require that additional Serial Bus-dependent registers shall be implemented, as summarized by the table below.

Offset	Register name	Description
200 ₁₆	CYCLE_TIME	24.576 MHz clock required for isochronous operation
204 ₁₆	BUS_TIME	System time in seconds
21C ₁₆	BUS_MANAGER_ID	Contains the <i>node_ID</i> of the bus manager, if one is present
220 ₁₆	BANDWIDTH_AVAILABLE	Known location for Serial Bus isochronous bandwidth allocation
224 ₁₆ – 228 ₁₆	CHANNELS_AVAILABLE	Known location for Serial Bus isochronous channel allocation

6.3 MANAGEMENT_AGENT register

The MANAGEMENT_AGENT register permits the initiator to signal the address of a management ORB to the target. This register shall support 8-byte block read and block write requests whose *destination_offset* is equal to the address of the MANAGEMENT_AGENT register and shall reject quadlet write requests and all other block read and block write requests. The format of this register is illustrated below.

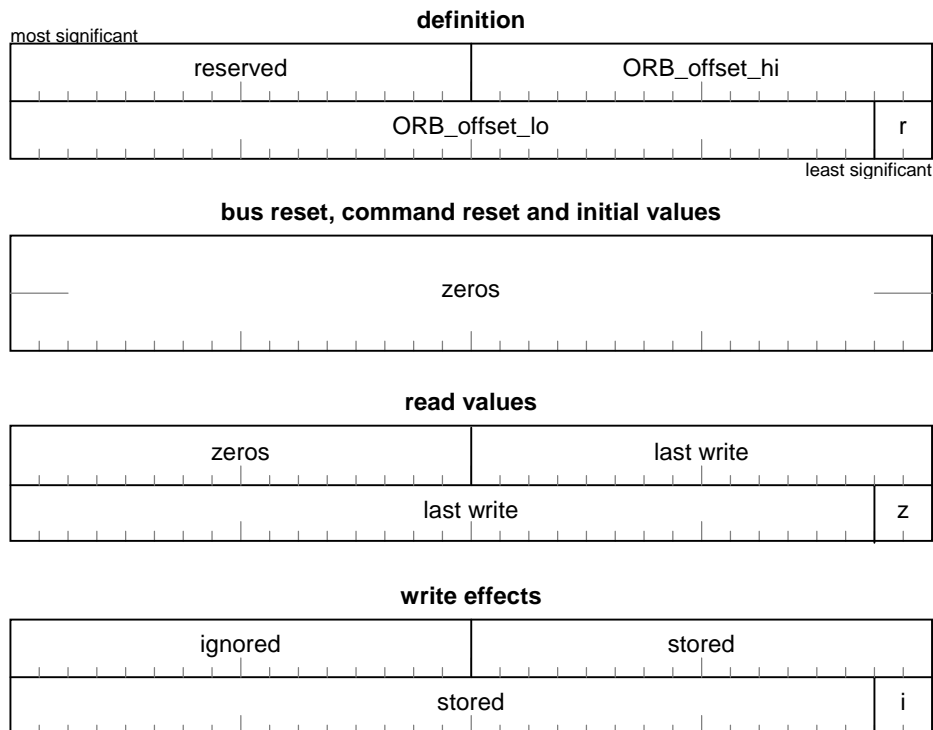


Figure 34 – MANAGEMENT_AGENT format

The *ORB_offset_hi* and *ORB_offset_lo* fields together form an *ORB_offset* field from which a Serial Bus address is derived when the management ORB is fetched. The Serial Bus address shall be formed from

the concatenation of the 16-bit node ID of the initiator (available to the target as the *source_ID* field of the block write request that updated the register), the *ORB_offset* field and two least significant bits of zero.

An initiator may signal a request by means of an 8-byte block write transaction that specifies the address of the request. If the management agent is busy with another request, the block write shall be rejected with a response of *resp_conflict_error*. If the write transaction is successful, the management agent shall fetch the request specified by *ORB_offset* and execute it. Unsuccessful write transactions shall not affect the execution of any request(s) in progress.

Because IEEE Std 1394-1995 reserves a portion of initial units space for bus-dependent use, the MANAGEMENT_AGENT register shall be located at address FFFF F001 0000₁₆ or above within the node's 48-bit address range. The address of the management agent is specified by the *csr_offset* field in the Management_Agent entry in configuration ROM (see 7.4.7).

6.4 Command block and stream control agent registers

Unlike the management agent, which services a single request at a time, the command block and stream control agents manage linked lists from which they fetch requests. For this reason they are referred to as fetch agents. Each target fetch agent has a set of control and status registers that lie within the target's initial units space; the fetch agent CSR's shall be located at or above address FFFF F001 0000₁₆ within the node's 48-bit address range.

Although the location of each fetch agent's CSR's is not fixed, the relative relationship of the registers is fixed within a contiguous block of eight quadlets, as defined by the table below.

Relative offset	Name	Description
00 ₁₆	AGENT_STATE	Reports fetch agent state
04 ₁₆	AGENT_RESET	Resets fetch agent
08 ₁₆	ORB_POINTER	Address of ORB
10 ₁₆	DOORBELL	Signals fetch agent to refetch an address pointer
14 ₁₆	UNSOLICITED_STATUS_ENABLE	Acknowledges the initiator's receipt of unsolicited status
18 ₁₆ – 1C ₁₆		Reserved for future standardization

The base address of each set of fetch agent's CSR's is obtained from the *command_block_agent* field in the login response returned by the target as part of a successful login.

6.4.1 AGENT_STATE register

The AGENT_STATE register is a read-only register that provides information about the current condition of the fetch agent. The definition is given by Figure 35.




Figure 35 – AGENT_STATE format

The *st* field shall specify the current operational state of the fetch agent, as encoded by the values in the table below.

Value	Fetch agent state
0	RESET
1	ACTIVE
2	SUSPENDED
3	DEAD

6.4.2 AGENT_RESET register

The AGENT_RESET register permits an initiator to reset the operational state of a target fetch agent. The definition of this write-only register is given by Figure 36.




Figure 36 – AGENT_RESET format

A quadlet write of any value to this register shall cause all fetch agent CSR's to be reset to their initial values, after which the fetch agent shall transition to the reset state.

6.4.3 ORB_POINTER register

The ORB_POINTER register contains the address of an ORB in system memory. This register shall support 8-byte block read and block write requests whose *destination_offset* is equal to the address of the ORB_POINTER register and shall reject quadlet write requests and all other block read and block write requests. The definition is given by Figure 37.




Figure 37 – ORB_POINTER format

The *ORB_offset_hi* and *ORB_offset_lo* fields together form an *ORB_offset* field from which a Serial Bus address is derived when the ORB is fetched. The Serial Bus address shall be formed from the concatenation of the 16-bit node ID of the initiator (available to the target as a result of a login), the *ORB_offset* field and two least significant bits of zero.

The effects of a write transaction to the ORB_POINTER register are dependent upon the value of *st* in the AGENT_STATE register. If the target agent is in the DEAD state, writes to the ORB_POINTER register shall be ignored. If the target agent is in the RESET or SUSPENDED state, a write to this register shall cause the *ORB_offset* to be stored and the agent to transition to the ACTIVE state. If the target agent is in the ACTIVE state, a write to the ORB_POINTER register may cause unpredictable target behavior.

6.4.4 DOORBELL register

The DOORBELL register provides a means by which the initiator signals the target that a linked list of requests has been updated. The definition of this write-only register is given by Figure 38.




Figure 38 – DOORBELL format

A quadlet write of any value to this register shall cause the fetch agent's *doorbell* variable to be set to one.

6.4.5 UNSOLICITED_STATUS_ENABLE register

The UNSOLICITED_STATUS_ENABLE register provides a means by which the initiator may grant the target permission to store an unsolicited status block. The definition of this write-only register is given by Figure 39.




Figure 39 – UNSOLICITED_STATUS_ENABLE format

A quadlet write of any value to this register shall cause the fetch agent's *unsolicited status enabled* variable to be set to one. A successful login or create stream request shall zero the *unsolicited status enabled* variable. As described in 9.4, any time a target stores an unsolicited status block it shall zero the *unsolicited status enabled* variable for that login. Before the target may store a subsequent unsolicited status block it is necessary for the initiator to write to the UNSOLICITED_STATUS_ENABLE register.

7 Configuration ROM

All nodes that implement SBP-2 targets shall implement general format configuration ROM in accordance with ISO/IEC 13213:1994, IEEE Std 1394-1995 and this standard. General format configuration ROM is a self-descriptive structure as illustrated below. The bus information block and root directory are at fixed locations; all other directories and leaves are addressed by entries in their parent directory.

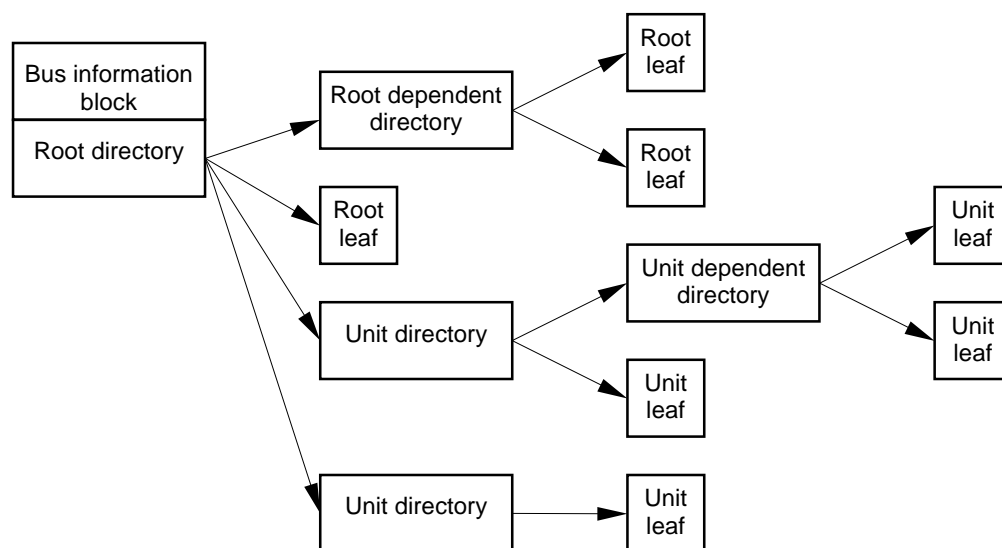


Figure 40 – Configuration ROM hierarchy

The figure above shows the potential of the general ROM format to accommodate a diversity of directory and leaf entries in a tree structure. In practice a target need implement only a portion of the entries shown above.

7.1 Power reset initialization

During the initialization process that follows a power reset a target may not be able to respond to Serial Bus request subactions addressed to parts of configuration ROM. When the target has insufficient information to make more than the first quadlet of configuration ROM accessible, it shall return a response value of zero for any read request addressed to FFFF F000 0400₁₆ or acknowledge the request subaction with *ack_tardy*, as specified by draft standard IEEE P1394a.

Targets shall complete initialization within five seconds of a power reset. Once power reset initialization completes, the target shall make all mandatory configuration ROM entries available. The target should not initiate a Serial Bus reset solely as a consequence of the completion of power reset initialization.

Optional configuration ROM information, such as textual descriptor leaves that identify the target vendor and model, may not be available when power reset initialization completes. The target may add this information to configuration ROM as it becomes available and may initiate a Serial Bus reset to alert other nodes to the changed configuration ROM. The target should initiate a Serial Bus reset if there is no expectation that other nodes would otherwise become aware of changed configuration ROM.

7.2 Bus information block

All targets shall implement a bus information block at a base address of FFFF F000 0404₁₆. For convenience of reference, the format of the bus information block defined by IEEE Std 1394-1995 is

reproduced below. The current version of the referenced standard or its supplements shall be consulted for the most recent information.

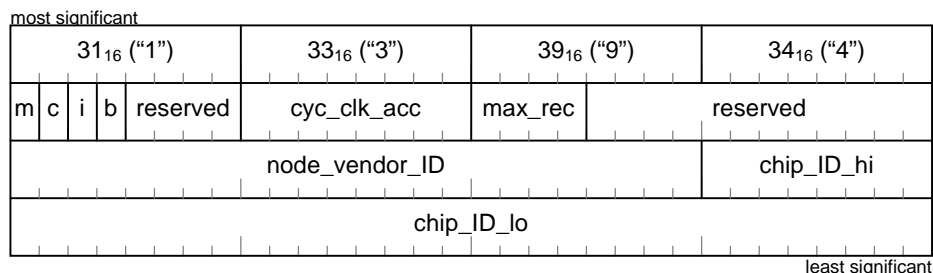


Figure 41 – Bus information block format

The first quadlet contains the string "1394" in ASCII characters.

The *irmc* bit (abbreviated as *m* in the figure above) shall be one if the node is isochronous resource manager capable; otherwise, the *irmc* value shall be zero.

The *cmc* bit (abbreviated as *c* in the figure above) shall be one if the node is cycle master capable; otherwise, this value shall be zero.

The *isc* bit (abbreviated as *i* in the figure above) shall be one if the node supports isochronous operations; otherwise, this value shall be zero.

The *bmc* bit (abbreviated as *b* in the figure above) shall be one if the node is bus manager capable; otherwise, this value shall be zero.

The *cyc_clk_acc* field specifies the node's cycle master clock accuracy in parts per million. If the *cmc* bit is one, the value in this field shall be between zero and 100. If the *cmc* bit is zero, this field shall be all ones.

The *max_rec* field defines the maximum payload size of a block write request addressed to the node. The range of the maximum payload size is from four bytes to 2048 bytes. A *max_rec* value of zero indicates that the maximum payload size is not specified. Otherwise, within the range of defined payload sizes, the maximum size is equal to $2^{max_rec + 1}$. The *max_rec* field does not place any limits on the maximum payload size in asynchronous data packets, either requests or responses, that the target may transmit.

The *node_vendor_ID* field shall be uniquely assigned by the IEEE/RAC, as specified by ISO/IEC 13213:1994. Unique identifiers for a company or organization may be obtained from:

Institute of Electrical and Electronic Engineers, Inc.
Registration Authority Committee
445 Hoes Lane
Piscataway, NJ 08855-1331

The *chip_ID_hi* and *chip_ID_lo* fields are concatenated to form a 40-bit chip ID value. The vendor specified by *node_vendor_ID* shall administer the chip ID values. When appended to the *node_vendor_ID* value, these shall form a unique 64-bit value called the EUI-64 (Extended Unique Identifier, 64 bits). The EUI-64 is also referred to as the node unique ID. Because physical addresses on Serial Bus may change after a bus reset, this unique identifier is the only reliable method of node identification.

7.3 Root directory

Configuration ROM for targets shall contain a root directory. The root directory immediately follows the bus information block and has a base address of FFFF F000 0414₁₆. The root directory shall contain Module_Vendor_ID and Node_Capabilities entries.

The root directory shall also contain at least one Unit_Directory entry that specifies the location of a unit directory whose format is specified by this standard.

7.3.1 Module_Vendor_ID entry

The Module_Vendor_ID entry is an immediate entry in the root directory that provides the company ID of the vendor that manufactured the module. Figure 42 shows the format of this entry.




Figure 42 – Module_Vendor_ID entry format

03₁₆ is the concatenation of *key_type* and *key_value* for the Module_Vendor_ID entry.

The IEEE/RAC uniquely assigns the *module_vendor_ID* to each module vendor, as specified by ISO/IEC 13213:1994. There is no requirement that the values of *module_vendor_ID* and *node_vendor_ID* be equal.

NOTE – A recommended convention to provide vendor identification in displayable form is to immediately follow the Module_Vendor_ID entry with a textual descriptor leaf entry. This associates an ASCII string with the module vendor. See ISO/IEC 13213:1994 for the specification of textual descriptor leaves; examples are given in Annex D.

7.3.2 Node_Capabilities entry

The Node_Capabilities entry is an immediate entry in the root directory that describes node capabilities. Figure 43 shows the format of this entry.




Figure 43 – Node_Capabilities entry format

0C₁₆ is the concatenation of *key_type* and *key_value* for the Node_Capabilities entry.

The *node_capabilities* field contains subfields, specified by ISO/IEC 13213:1994. Targets shall implement the SPLIT_TIMEOUT register, the 64-bit fixed addressing scheme, the STATE_CLEAR.*lost* bit and the STATE_CLEAR.*dreq* bit and indicate this by setting the *spt*, *64*, *fix*, *lst* and *drq* bits to one. If no other *node_capabilities* bits are one this results in a value of 00 83C0₁₆.

7.3.3 Unit_Directory entry

The Unit_Directory entry is a directory entry in the root directory that describes the location of a unit directory within configuration ROM. There may be more than one unit directory; each unit directory shall be located by a separate Unit_Directory entry. Figure 44 shows the format of this entry.



Figure 44 – Unit_Directory entry format

D1₁₆ is the concatenation of *key_type* and *key_value* for the Unit_Directory entry.

The *indirect_offset* field specifies the number of quadlets from the address of the Unit_Directory entry to the address of the unit directory within configuration ROM.

7.4 Unit directory

Configuration ROM for targets shall contain at least one unit directory in the format specified by this standard. The unit directory shall contain Unit_Spec_ID and Unit_SW_Version entries, as specified by ISO/IEC 13213:1994, and a Management_Agent entry, as specified by this standard.

Targets shall implement at least one logical unit, logical unit zero. Additional logical units may be implemented. A logical unit is described by entries in the unit directory or by entries in a logical unit directory dependent upon the unit directory or by entries taken in combination from both places. The properties of logical units are established by Command_Set_Spec_ID, Command_Set, Command_Set_Revision and Logical_Unit_Characteristics entries; an instance of a specific logical unit is established by a Logical_Unit_Number entry.

The unit directory may also contain a Unit_Unique_ID entry.

7.4.1 Unit_Spec_ID entry

The Unit_Spec_ID entry is an immediate entry in the unit directory that specifies the organization responsible for the architectural definition of the target. Figure 45 shows the format of this entry.



Figure 45 – Unit_Spec_ID entry format

12₁₆ is the concatenation of *key_type* and *key_value* for the Unit_Spec_ID entry.

00 609E₁₆ is the *unit_spec_ID* obtained by NCITS from the IEEE/RAC. The value indicates that the NCITS Secretariat is responsible for the software interface definition.

7.4.2 Unit_SW_Version entry

The Unit_SW_Version entry is an immediate entry in the unit directory that, in combination with the *unit_sw_version*, specifies the software interface of the target. Figure 46 shows the format of this entry.



Figure 46 – Unit_SW_Version entry format

13₁₆ is the concatenation of *key_type* and *key_value* for the Unit_SW_Version entry.

01 0483₁₆ is the *unit_sw_version* value that indicates that the target conforms to this standard.

7.4.3 Command_Set_Spec_ID entry

The Command_Set_Spec_ID entry is an immediate entry that, when present in the unit directory, specifies the organization responsible for the command set definition for the target. Figure 47 shows the format of this entry.




Figure 47 – Command_Set_Spec_ID entry format

38₁₆ is the concatenation of *key_type* and *key_value* for the Command_Set_Spec_ID entry.

The *command_set_spec_ID* is an organizationally unique identifier obtained from the IEEE/RAC. The organization to which this 24-bit identifier has been granted is responsible for the definition of the command set implemented by the target.

7.4.4 Command_Set entry

The Command_Set entry is an immediate entry that, when present in the unit directory, in combination with the *command_set_spec_ID* specifies the command set implemented by the target. Figure 48 shows the format of this entry.



Figure 48 – Command_Set entry format

39₁₆ is the concatenation of *key_type* and *key_value* for the Command_Set entry.

The value of *command_set* shall be specified by the owner of *command_set_spec_ID*.

7.4.5 Command_Set_Revision entry

The Command_Set_Revision entry is an immediate entry that, when present in the unit directory, specifies the revision level of the command set implemented by the target. Figure 49 shows the format of this entry.




Figure 49 – Command_Set_Revision entry format

3B₁₆ is the concatenation of *key_type* and *key_value* for the Command_Set_Revision entry.

The value of *command_set_revision* shall be specified by the owner of *command_set_spec_ID*.

7.4.6 Firmware_Revision entry

The Firmware_Revision entry is an immediate entry that, when present in the unit directory, specifies the firmware revision level implemented by the target. Figure 50 shows the format of this entry.



Figure 50 – Firmware_Revision entry format

3C₁₆ is the concatenation of *key_type* and *key_value* for the Firmware_Revision entry.

The value of *firmware_revision* shall be specified by the manufacturer of the target.

7.4.7 Management_Agent entry

The Management_Agent entry is an immediate entry in the unit directory that specifies the base address of the target's MANAGEMENT_AGENT register. Figure 51 shows the format of this entry.



Figure 51 – Management_Agent entry format

54₁₆ is the concatenation of *key_type* and *key_value* for the Management_Agent entry.

The *csr_offset* field shall contain the offset, in quadlets, from the base address of initial register space, FFFF F000 0000₁₆, to the base address of the MANAGEMENT_AGENT register for the target. All target CSR's shall be located at or above address FFFF F001 0000₁₆; therefore the value of *csr_offset* shall not be less than 4000₁₆.

NOTE – If a device implements additional control and status registers that are dependent upon the device class, it is recommended that these registers be placed at one of two locations within the device's address space. If the additional register(s) pertain to a logical unit, the recommended locations are at offset 20₁₆ and above following the base address of the logical unit's command block agent registers. Additional register(s) that are associated with the device, and not a particular logical unit, may be located immediately after the MANAGEMENT_AGENT register. If this convention is followed, there is no necessity for additional configuration ROM entries to describe the location of device-dependent registers.

7.4.8 Logical_Unit_Characteristics entry

The Logical_Unit_Characteristics entry is an immediate entry that, when present in the unit directory, specifies characteristics of the target implementation. Figure 52 shows the format of this entry.




Figure 52 – Logical_Unit_Characteristics entry format

3A₁₆ is the concatenation of *key_type* and *key_value* for the Logical_Unit_Characteristics entry.

The *q* bit shall specify the task management (queuing) model implemented by the target. If *q* is zero, the target implements the basic task management model defined by this standard in 10.2. When *q* is one, the task management model is dependent upon the command set specified by the Command_Set_Spec_ID and Command_Set entries.

The *ordered* bit (abbreviated as *o* in the figure above) specifies the manner in which the target executes tasks signaled to the normal command block agent. If the target executes and reports completion status without any ordering constraints, the *ordered* bit shall be zero. Otherwise, if the target both executes all tasks in order and reports their completion status in the same order, the *ordered* bit shall be one.

The *isochronous* bit (abbreviated as *i* in the figure above) specifies whether or not the target supports isochronous operations. When *isochronous* is one, create stream requests and stream command block requests shall be supported; stream control requests may be supported. If the *isochronous* bit is one, the *irmc*, *cmc* and *isc* bits in the bus information block shall also be one, as described in 7.2.

The *mgt_ORB_timeout* field shall specify, in units of 500 milliseconds, the maximum time an initiator shall allow for a target to store a status block in response to a management ORB. The time-out commences when the initiator receives either *ack_complete* or *resp_complete* from the target in response to the block write of the management ORB address to the MANAGEMENT_AGENT register.

The *ORB_size* field shall specify, in quadlets, the fetch size used by the target to obtain ORB's from initiator memory. The initiator shall allocate, on a quadlet aligned boundary, at least this much memory for each ORB signaled to the target.

7.4.9 Logical_Unit_Directory entry

The Logical_Unit_Directory entry is an optional directory entry in the unit directory that describes the location of the logical unit directory within configuration ROM. Figure 53 shows the format of this entry.



Figure 53 – Logical_Unit_Directory entry format

D4₁₆ is the concatenation of *key_type* and *key_value* for the Logical_Unit_Directory entry.

The *indirect_offset* field specifies the number of quadlets from the address of the Logical_Unit_Directory entry to the address of the logical unit directory within configuration ROM.

7.4.10 Logical_Unit_Number entry

The Logical_Unit_Number entry is an immediate entry that, when present in the unit directory, specifies the peripheral device type and logical unit number of a logical unit implemented by the target. Figure 54 shows the format of this entry.




Figure 54 – Logical_Unit_Number entry format

14₁₆ is the concatenation of *key_type* and *key_value* for the Logical_Unit_Number entry.

The *device_type* field indicates the peripheral device type implemented by the logical unit. This field shall contain a value specified by the table below.

Value	Peripheral device type
0 – 1E ₁₆	The meaning of <i>device_type</i> is command set-dependent
1F ₁₆	Unknown device type; command set-dependent means are necessary to determine the peripheral device type

The *lun* field shall identify the logical unit to which the information in the Logical_Unit_Number entry applies.

7.4.11 Unit_Unique_ID entry

The Unit_Unique_ID entry is an optional leaf entry in the unit directory that describes the location of the unit unique ID leaf within configuration ROM. If a vendor implements a device with multiple Serial Bus access paths, *i.e.*, multiple links to Serial Bus each of which receives a distinct *node_ID* as the result of Serial Bus initialization or bus enumeration, the Unit_Unique_ID entry shall be implemented. Figure 55 shows the format of this entry.



Figure 55 – Unit_Unique_ID entry format

8D₁₆ is the concatenation of *key_type* and *key_value* for the Unit_Unique_ID entry.

The *indirect_offset* field specifies the number of quadlets from the address of the Unit_Unique_ID entry to the address of the unit unique ID leaf within configuration ROM.

7.5 Logical unit directory

The logical unit directory provides one of two methods by which a logical unit implemented by the target may be described (the other is a Logical_Unit_Number entry in the unit directory, already described in 7.4.10). The entries permitted in a logical unit directory are summarized in the table below.

Entry	Required
Command_Set_Spec_ID	Optional
Command_Set	Optional
Command_Set_Revision	Optional
Logical_Unit_Characteristics	Optional
Logical_Unit_Number	Mandatory

7.5.1 Command_Set_Spec_ID entry

The Command_Set_Spec_ID entry is an immediate entry that, when present in a logical unit directory, specifies the organization responsible for the command set definition for the logical unit. If there is no Command_Set_Spec_ID entry in the logical unit directory, the Command_Set_Spec_ID entry in the unit directory shall apply; otherwise the entry in the logical unit directory shall take precedence. Figure 47 shows the format of this entry; the fields are defined in 7.4.3.

7.5.2 Command_Set entry

The Command_Set entry is an immediate entry that, when present in a logical unit directory and in combination with the *command_set_spec_ID*, specifies the command set implemented by the logical unit. If there is no Command_Set entry in the logical unit directory, the Command_Set entry in the unit directory shall apply; otherwise the entry in the logical unit directory shall take precedence. Figure 48 shows the format of this entry; the fields are defined in 7.4.4.

7.5.3 Command_Set_Revision entry

The Command_Set_Revision entry is an immediate entry that, when present in a logical unit directory, specifies the revision level of the command set implemented by the logical unit. If there is no Command_Set_Revision entry in the logical unit directory, the Command_Set_Revision entry in the unit directory shall apply; otherwise the entry in the logical unit directory shall take precedence. Figure 49 shows the format of this entry; the fields are defined in 7.4.5.

7.5.4 Logical_Unit_Characteristics entry

The Logical_Unit_Characteristics entry is an immediate entry that, when present in a logical unit directory, specifies characteristics of the logical unit implementation. If there is no Logical_Unit_Characteristics entry in the logical unit directory, the Logical_Unit_Characteristics entry in the unit directory shall apply; otherwise the entry in the logical unit directory shall take precedence. Figure 52 shows the format of this entry; the fields are defined in 7.4.8.

7.5.5 Logical_Unit_Number entry

The Logical_Unit_Number entry is an immediate entry in a logical unit directory that specifies peripheral device type and logical unit number of the logical unit implementation. Figure 54 shows the format of this entry; the fields are defined in 7.4.10.

7.6 Unit unique ID leaf

Although the node unique ID (EUI-64) present in the bus information block is sufficient to uniquely identify nodes attached to Serial Bus, it is insufficient to identify a target when a vendor implements a device with multiple Serial Bus node connections. In this case initiator software requires information by which a particular target may be uniquely identified, regardless of the Serial Bus access path used. The figure below shows the format of the unit unique ID leaf.




Figure 56 – Unit unique ID leaf format

The first quadlet of the unit unique leaf shall contain the number of following quadlets in the leaf and a CRC calculated for those quadlets, as specified by ISO/IEC 13213:1994.

The *unit_vendor_ID* value shall be uniquely assigned by the IEEE/RAC, as specified by ISO/IEC 13213:1994.

The *unit_ID_hi* and *unit_ID_lo* fields are concatenated to form a 40-bit unit ID value. The vendor specified by *unit_vendor_ID* shall administer the unit ID values. When appended to the *unit_vendor_ID* value, these shall form a unique 64-bit value referred to as the unit unique ID.

As a consequence of the implementation of multiple Serial Bus nodes for the same unit, there is configuration ROM accessible for each node. Parts of these configuration ROM's shall differ from each other, *e.g.*, the node unique ID in the bus information block, but the unit unique ID shall be the same regardless of which node is used to access the information.

8 Access

Before an initiator may signal commands or other requests to a target, access privileges shall first be granted by the target. The criteria for the grant of access may include resource availability or other target requirements. This section specifies the target facilities that support access control and the methods by which an initiator requests access to a target and eventually relinquishes access when it is no longer required.

8.1 Access protocols

Targets shall implement a logical unit reservation protocol that by itself supports neither persistent reservations nor passwords; it is a simple mechanism that can be used to guarantee single initiator access to the logical unit and to preserve that initiator's access rights across a Serial Bus reset.

In order to support the logical unit reservation protocol, a target shall implement resources to manage one or more logins from initiators. These resources are described below and are used in the specification of target actions in response to login requests signaled by an initiator to the target's management agent:

- The target implements a set of one or more *login_descriptors* that are used to hold context for logins. The context of a login stored in a *login_descriptor* consists of the *lun*, the *login_owner_ID*, the *login_owner_EUI_64*, the *status_FIFO* address, an *exclusive* variable, the base addresses of the fetch agent CSR's returned to the initiator in the *login_response* data and the *login_ID* used by the initiator to identify the login.
- The *login_owner_ID* is the 16-bit node ID of the current owner of a login. Upon either a Serial Bus reset or a power reset, the *login_owner_ID* for all *login_descriptors* shall be reset to all ones. The target shall use the *login_owner_ID* to qualify all write requests addressed to the *login_descriptor* fetch agent CSR's.
- The *login_owner_EUI_64* is the unique 64-bit identifier of the current owner of a login. Upon a power reset, the *login_owner_EUI_64* for all *login_descriptors* shall be reset to all ones. Upon a Serial Bus reset, the *login_owner_EUI_64* persists for two seconds and shall then be reset to all ones unless it has been reestablished.

A *login_descriptor* is considered free if both its *login_owner_ID* and *login_owner_EUI_64* are all ones. The resources of this *login_descriptor* may be allocated to any initiator that successfully completes a login or a create stream request. If a *login_descriptor's* *login_owner_ID* is all ones but its *login_owner_EUI_64* holds a valid EUI-64, the *login_descriptor* is reserved—the initiator identified by *login_owner_EUI_64* may reestablish the login. Active *login_descriptors* are those whose *login_owner_ID* and *login_owner_EUI_64* are both valid; the initiator that owns the login may signal requests to the fetch agent(s) associated with the *login_descriptor*.

8.2 Login requests

The clauses that follow describe the use of the login ORB's defined in 5.1.4.

8.2.1 Login

Before an initiator may signal any other requests to a target it shall first perform a login. The login request, whose format is specified in 5.1.4.1, shall be signaled to the target's MANAGEMENT_AGENT register by means of an eight-byte block write transaction that specifies the Serial Bus address of the login request. The address of the management agent shall be obtained from configuration ROM.

NOTE – The speed at which the block write request to the MANAGEMENT_AGENT register is received shall determine the speed used by the target for all subsequent requests to read the initiator's configuration ROM,

fetch ORB's from initiator memory or store status at the initiator's *status_FIFO*. Command block ORB's separately specify the speed for requests addressed to the data buffer or page table.

The login ORB shall specify the *lun* of the logical unit for which the initiator desires access.

The target shall perform the following to validate a login request:

- a) The target shall read the initiator's unique ID, EUI-64, from the bus information block by means of two quadlet read transactions. The *source_ID* from the write transaction used to signal the login ORB to the target's MANAGEMENT_AGENT register shall be used as the *destination_ID* in the quadlet read transactions;
- b) The target shall determine whether or not the initiator already owns a login by comparing the EUI-64 just obtained against the *login_owner_EUI_64* for all *login_descriptors*. If the initiator is currently logged-in to the same logical unit, the login request shall be rejected;
- c) If the *exclusive* bit is set in the login ORB, the target shall reject the login request if there are any active *login_descriptors* for the logical unit;
- d) If an active *login_descriptor* with the *exclusive* attribute exists for the *lun* specified in the login ORB, the target shall reject the login request; and
- e) The target shall determine if a free *login_descriptor* is available. If a *login_descriptor* is free, the initiator's *source_ID* is stored in *login_owner_ID*, the initiator's EUI-64 is stored in *login_owner_EUI_64*, the *lun* and *status_FIFO* fields from the login ORB are stored in the *login_descriptor*, the *exclusive* variable in the *login_descriptor* is set to the value of the *exclusive* bit from the login ORB and the addresses of the fetch agent(s) are stored in the *login_descriptor*. Lastly the target assigns a unique *login_ID* to this login and stores it in the *login_descriptor*.

If the target is able to satisfy the login request, it shall return a login response as specified in 5.1.4.1. A critical component of a login response returned to the initiator is the base address of the target agent that the initiator shall use to signal any subsequent requests to the target for the indicated *login_ID*.

8.2.2 Create stream

An isochronous stream may be created for an initiator only after completion of the login process just described. The initiator shall supply a *login_ID* previously obtained as the result of a successful login as well as other information in the create stream request that characterizes the isochronous operations to be performed.

The isochronous information consists of three items:

- whether the target is to function as a talker or a listener;
- the maximum number of channels that may be simultaneously enabled; and
- the aggregate maximum isochronous payload for all channels to be transferred between Serial Bus and the device medium in a single isochronous cycle.

The initiator shall specify which role, talker or listener, the target shall assume for this stream. The maximum number of channels is required in order for the target to allocate sufficient resources.

The aggregate maximum isochronous payload is the worst-case amount of data the target may have to transfer to or from Serial Bus and from or to the medium in an isochronous cycle. Implementation-dependent constraints may limit the performance of the target, which requires this information in order to determine if the login may be accepted.

NOTE – Upon playback (when the target is a talker), the aggregate maximum isochronous payload shall reflect the total of all channels recorded on the medium—not just the aggregation of payload(s) for the channels to be

transmitted on Serial Bus. This is essential since the target reads all of the data from the medium even though the channel mask may select a small subset for playback.

The target shall perform the following to validate a create stream request:

- a) The target shall validate the *login_ID* supplied in the create stream ORB by comparing the *destination_ID* in the read request(s) used to fetch the ORB with the *source_ID* retained when *login_ID* was assigned to the initiator. If the node ID's do not match, the *login_ID* is invalid.
- b) If the *login_ID* is valid, the target shall determine if a free *login_descriptor* is available. If a *login_descriptor* is free, the initiator's *source_ID* is stored in *login_owner_ID*, the initiator's EUI-64 is stored in *login_owner_EUI_64*, the *lun* from the *login_descriptor* is copied to the *login_descriptor* for the create stream request and the addresses of the fetch agent(s) are also stored in the *login_descriptor*. Lastly the target assigns a unique *login_ID* to this login and stores it in the *login_descriptor*.

In addition to the addresses of the stream command block and stream control fetch agents, the target shall also specify in the *login_response* data the minimum transfer length that the initiator should specify in the *stream_length* field of any stream command block request signaled to the target.

8.3 Reconnection

Upon a Serial Bus reset, the target shall abort all task sets for all command block agents created as the result of login request(s). Task sets associated with isochronous streams shall not be aborted. Both the stream command block and stream control requests shall continue to be executed by the target but the return of status shall be deferred until a successful reconnection.

For one second subsequent to a bus reset the target shall retain sufficient information to permit an initiator to reconnect its login ID and associated stream ID's. After one second the target shall perform an implicit logout for all login ID's and stream ID's that have not been successfully reconnected to their original initiator(s).

NOTE – The basis of the one second time-out is to permit initiators to reallocate isochronous channels and bandwidth and to reestablish isochronous connections. The time-out commences when the target observes the first subaction gap subsequent to a bus reset. If a bus reset occurs before the time-out expires, the timer is zeroed then restarted upon detection of a subaction gap.

The target shall perform the following to validate a reconnect request:

- a) The target shall read the initiator's unique ID, EUI-64, from the bus information block by means of two quadlet read transactions. The *source_ID* from the write transaction used to signal the reconnect ORB to the target's MANAGEMENT_AGENT register shall be used as the *destination_ID* in the quadlet read transactions;
- b) The target shall determine whether or not the *login_ID* is valid by comparing the just obtained EUI-64 against the *login_owner_EUI_64* for the *login_descriptor* identified by *login_ID*;
- c) If the *login_ID* is valid, the target shall store the initiator's *source_ID* in *login_owner_ID* for the referenced *login_descriptor* and for all stream descriptors associated with the same initiator; and
- d) Fetch agents for stream command block and stream control requests for the reconnected initiator may resume; status for completed ORB's that had not been stored in the initiator's *status_FIFO* (because the initiator's *source_ID* had been invalidated by the bus reset) may also be stored.

No *login_response* data is stored for a reconnect request; the completion status is indicated by the status block stored at the *status_FIFO* address.

8.4 Logout

When an initiator no longer requires access to a target's resources, it shall signal a logout request to the management agent. The login or stream resources to be released shall be identified by *login_ID* in the logout ORB. A target shall reject a logout request if *login_ID* does not match that of any active *login_descriptor* or if the *source_ID* of the write request used to signal the logout ORB to the MANAGEMENT_AGENT register is not equal to the *source_ID* of the matching *login_descriptor*. If any tasks or stream control ORB's are active at the time of the logout request, they shall be aborted in the same fashion as if the task set had been aborted. Upon successful completion of a logout request, all resources allocated to the initiator are free once again and may be used by the target to satisfy subsequent login requests.

9 Command execution

This section describes the procedures used by an initiator to request command execution by a target. As described in the model, requests are specified by data structures in system memory that are subsequently fetched by the target. While a target executes a request, it is responsible for any data transfer associated with the request. Once a request completes, successfully or in error, a status block is stored in system memory by the target. The data structures are defined in section 5; the initiator procedures for the use of these request and status blocks are described in the clauses that follow

9.1 Requests and request lists

Management requests (which include login and logout requests) are signaled to the target agent by means of a Serial Bus block write transaction that specifies the address of the management ORB. The management agent becomes busy while executing a request and refuses subsequent Serial Bus transactions until the current request is completed. The management agent does not require any initialization procedures.

The other target agents, command block and stream control, are characterized as fetch agents since they manage linked lists of requests in system memory and are responsible to fetch the ORB's. For normal command block, stream command block and stream control ORB's, the initiator produces requests and the target consumes them. These processes are asynchronous and independent of each other. Target efficiency is improved if the target can be kept occupied with an ample working set of requests. To this end, the initiator is permitted to arrange ORB's in linked lists and to dynamically append new requests to the lists while the target remains active.

Each normal command block, stream command block or stream control ORB contains an address pointer, *next_ORB*, which shall either be null or point to another ORB. A linked list of ORB's, previously illustrated by Figure 6, implicitly orders the ORB's—the fact that the ORB's are in order permits the target to execute them in order (or not) according to its device-dependent characteristics.

The target is responsible to fetch ORB's from system memory, as described in more detail in 9.1.4. The remainder of this clause describes what the initiator shall do to:

- initialize a target fetch agent;
- dynamically append new requests to an active list and notify a target fetch agent of the new requests; and
- notify a target fetch agent of a single new request.

9.1.1 Fetch agent initialization (informative)

After successful completion of a login procedure and the return of the base address of the fetch agent CSR's, the initiator may initialize the fetch agent as follows:

- a) The initiator allocates space for a dummy ORB and initializes it *per* the format described in 5.1.1. Although only the *next_ORB* field, *notify* bit and the *rq_fmt* field are significant within a dummy ORB, the initiator allocates at least the minimum ORB size specified by the target's configuration ROM. The initiator sets the *next_ORB* field to the null pointer value;
- b) The initiator resets the target fetch agent by a quadlet write to the fetch agent's AGENT_RESET register;

- c) The initiator writes the address of the dummy ORB to the fetch agent's ORB_POINTER register by means of an 8-byte block write request. In the example in Figure 57, this is the value 0000 0000 8004 00C0₁₆. This causes the fetch agent to transition to the ACTIVE state.

The figure below illustrates the result of these actions:




Figure 57 – Fetch agent initialization with a dummy ORB

When the fetch agent transitions to the active state as a result of the write to the ORB_POINTER register, it uses the value to fetch the dummy ORB (as target resources permit). The dummy ORB, by definition, completes immediately and the target fetch agent stores a status block for the request. However, the null pointer in the *next_ORB* field of the dummy ORB causes the fetch agent to transition to the suspended state. The ORB_POINTER register still points to the dummy ORB and the initiator may subsequently append additional requests, as described in 9.1.2.

9.1.2 Dynamic appends to request lists (informative)

Once a target fetch agent has been initialized and made active as described above, it is possible for the initiator to append new requests to the linked list while the fetch agent remains active. Assume that the initiator intends to add three new requests previously illustrated by Figure 6.

An initiator may append new requests to an active request list as follows:

- The initiator constructs a linked list of ORB's in system memory, as illustrated in the example. The *next_ORB* field of the last ORB contains a null pointer. The *next_ORB* fields of all other ORB's contain a valid pointer to a subsequent ORB;
- The initiator updates the *next_ORB* field of what had been the last ORB, in this example the dummy ORB in Figure 57, with the address of the first request in the new request list, in this example 0000 0000 8000 0000₁₆; and
- Lastly, the initiator transmits a quadlet write request, with any data value, to the fetch agent's DOORBELL register.

The final step informs the target that address pointers in the request list have been updated by the initiator. If the target fetch agent had not encountered a null pointer, the activation of the doorbell is redundant. However, if the target fetch agent is already suspended at the time *next_ORB* is updated, the activation of the doorbell is essential to reactivate the fetch agent. In this latter case, it is necessary for the target fetch agent to refetch all or part of an ORB from system memory in order to ascertain if a previously null pointer contains a valid address of an ORB.

NOTE – If the initiator has knowledge that the fetch agent is in the suspended state, the algorithm described above may be modified to write the address of the new ORB to the ORB_POINTER register in place of the write to the DOORBELL register. This has the virtue of avoiding a refetch of the *next_ORB* field from the ORB at which the fetch agent is suspended, but would produce unpredictable results if the fetch agent were not in the suspended state.

9.1.3 Fetch agent use by the BIOS (informative)

The BIOS, or any similar initiator application that executes in a single-threaded environment, has little need of the target fetch agent's capabilities to manage multiple outstanding requests. The BIOS may take advantage of this and use a simpler procedure than that described in 9.1.2 to signal requests to the target. Subsequent to initialization of the target fetch agent by means of a write to the AGENT_RESET register, the BIOS may signal one request at a time to the target as follows:

- a) The BIOS allocates space for the request in an ORB and initializes it according to the ORB format. The *next_ORB* field contains a null pointer;
- b) The BIOS signals the request to the target agent by writing the address of the ORB to the ORB_POINTER register in an 8-byte block write transaction. This causes the target agent to transition to the ACTIVE state and to execute the request; and
- c) Subsequent to the return of a status block to the *status_FIFO* address specified when the login was performed, the BIOS may signal additional requests by repeating this procedure.

The performance improvements yielded by the above procedure (which are accomplished by the elimination of a read transaction to fetch an ORB) are minor; the principal benefit to the BIOS is code simplification.

9.1.4 Fetch agent state machine

The operations of a target fetch agent are specified by the figure below. The state of a fetch agent is visible in the context displayed by the AGENT_STATE and ORB_POINTER registers described in 6.4. The state machine diagram and accompanying text explicitly specify the conditions for transition from one state to another and the actions taken within states.

The target shall qualify all writes to fetch agent CSR's by the *source_ID* of the currently logged-in initiator. A write to a fetch agent CSR by any other Serial Bus node shall be rejected by the target by one of the following methods:

- an acknowledgment of *ack_type_error*;
- an acknowledgment of *ack_complete* (although the write is ignored); or
- an acknowledgment of *ack_pending*. When the target subsequently responds, the response code shall be *resp_type_error*.

The recommended target action is to indicate a type error, either by an acknowledgment of *ack_type_error* or an acknowledgment of *ack_pending* followed by *resp_type_error*.

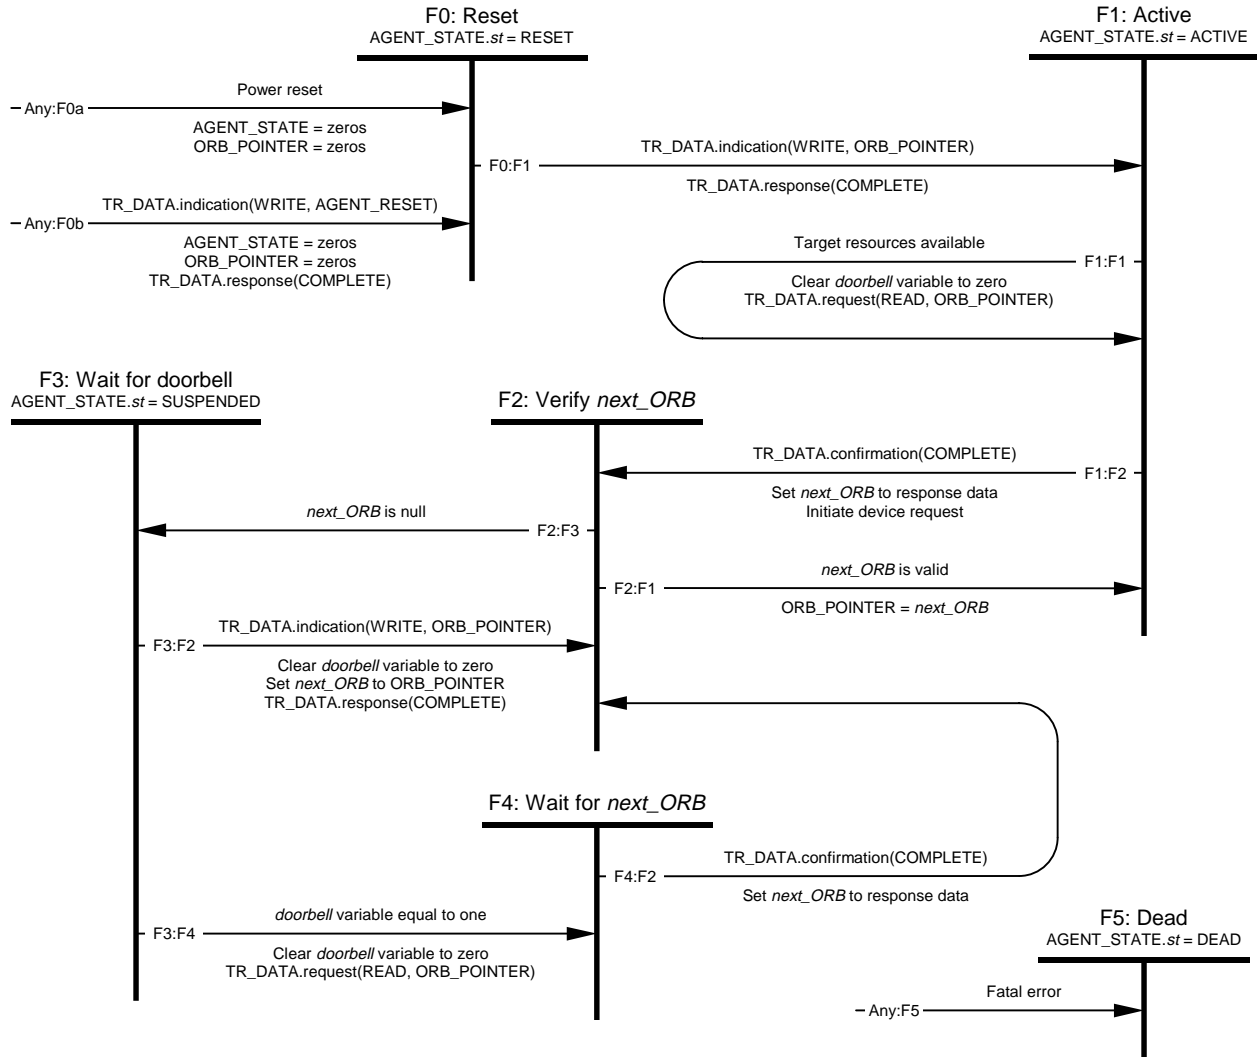


Figure 58 – Fetch agent state machine

Transition Any:F0a. A power reset shall cause the fetch agent to transition to the RESET state from any other state. The AGENT_STATE and ORB_POINTER registers (that control and make visible the operations of the fetch agent) shall be reset to zeros.

Transition Any:F0b. A quadlet write request by the initiator to the AGENT_RESET register shall cause the fetch agent to transition to state F0 from any other state. The fetch agent shall zero the AGENT_STATE and ORB_POINTER registers before the transition to state F0.

State F0: Reset. Upon entry to this state, the *st* field in the AGENT_STATE register shall be set to RESET. The fetch agent is inactive and available to be initialized by an initiator.

Transition F0:F1. An 8-byte block write of a valid *ORB_offset* to the ORB_POINTER register shall update the register and cause the fetch agent to transition to state F1. The target shall confirm the block write request with a response of COMPLETE.

NOTE – When the fetch agent is reset, it is not necessary to write to the DOORBELL register when a transition is made to the ACTIVE state.

State F1: Active. Upon entry to state F1, the *st* field in the AGENT_STATE register shall be set to ACTIVE. In this state, the fetch agent may use the address information in the ORB_POINTER register to fetch ORB's from the initiator as resources permit.

Transition F1:F1. The availability of target resources is an implementation-dependent decision. Typically, the resources might be space in device memory to hold an image of the ORB while the command is scheduled for execution and subsequently completed. In any case, the fetch agent clears the *doorbell* variable to zero and then issues a block read request to obtain the ORB from system memory.

Transition F1:F2. Subsequent to a block read request, issued as described above, the fetch agent may accept a block read response that contains the desired ORB. If a read response is received whose *source_ID*, *destination_ID* and *tl* fields match the *destination_ID*, *source_ID* and *tl* fields, respectively, of the read request, the fetch agent shall make the ORB available to the device server for execution and shall copy the *next_ORB* field from the response data to the *next_ORB* variable before making the transition to state F2. The target shall not initiate execution of the command contained within the ORB until these actions are complete.

State F2: Verify *next_ORB*. The *next_ORB* variable contains information about a subsequent ORB that may be linked in order after the one just fetched. As described in 5.1, the *next_ORB* pointer encodes the address of the next ORB. The actions of this state determine whether or not the *next_ORB* pointer is null.

Transition F2:F3. The fetch agent shall transition to a suspended state, F3, if *next_ORB* contains a null pointer. A null pointer is defined in 5.1 and exists if the most significant bit of the variable is one.

Transition F2:F1. If the *next_ORB* variable does not indicate a null pointer the fetch agent shall update the ORB_POINTER register with the value of *next_ORB*.

State F3: Wait for doorbell. The fetch agent is suspended; the ORB_POINTER register contains the address of the ORB that, at the time state F3 was entered contained a null pointer.

Transition F3:F2. If an indication of a write to the ORB_POINTER register is received, the fetch agent shall clear the *doorbell* variable to zero, set the *next_ORB* variable to the value of the ORB_POINTER register and then confirm the write transaction with a response of COMPLETE. After the confirmation, the fetch agent shall transition to state F2, from which state an immediate F2:F1 transition follows.

Transition F3:F4. Whenever the *doorbell* variable is equal to one, the fetch agent shall clear the *doorbell* variable to zero and then issue a read request to obtain a fresh copy of the *next_ORB* field from the ORB whose address is contained in the ORB_POINTER register and shall transition to state F4. The *doorbell* variable is set to one as the result of a quadlet write request of any value to the DOORBELL register, whether the write request is received in this or any other state.

NOTE – The fetch agent may issue either an 8-byte block read request (to fetch just the *next_ORB* field) or it may reread the entire ORB. The initiator shall insure that system memory occupied by the ORB remains accessible, as described in 9.3.

State F4: Wait for *next_ORB*. The fetch agent is suspended and awaiting a read response for a block read directed to the address contained in the ORB_POINTER register.

Transition F4:F2. Subsequent to a block read request, issued as described above, the fetch agent may accept a block read response that contains the *next_ORB* data. If a read response is received whose *source_ID*, *destination_ID* and *tl* fields match the *destination_ID*, *source_ID* and *tl* fields, respectively, of the read request, the fetch agent shall copy the *next_ORB* field from the response data to the *next_ORB* variable before making the transition to state F2 in order to verify the *next_ORB* variable. If, as expected, *next_ORB* is not null an immediate F2:F1 transition follows.

Transition Any:F5. Upon the detection of any fatal error, the fetch agent shall transition to state F5. Examples of fatal errors include, but are not limited to:

- the failure of the addressed node to acknowledge a read request;
- the failure of the addressed node to respond to a read request (split time-out);
- a busy condition at the addressed node that exceeds the target's busy retry limit;
- a data CRC error in a response.

Some of these errors may be recoverable if retried by the target.

State F5: Dead. The dead state is a unique state that preserves fetch agent information in the AGENT_STATE and ORB_POINTER registers. All writes to these registers shall have no effect while in state F5.

9.2 Data transfer

The transfer of data associated with a command is the responsibility of the target. The target shall use Serial Bus read transactions to fetch data from system memory and Serial Bus write transactions to store data in system memory.

The total transfer length may be larger than the maximum data payload that can be accommodated in a single transaction. The target is responsible to manage the size and number of read or write transactions to transfer all the requested data. The target may choose any appropriate size for these data transfer transactions, subject to constraints specified by the ORB.

The target shall observe alignment requirements specified by the *page_table_present* bit and the *page_size* field. If *page_table_present* is one, the target shall observe alignment boundaries that occur every 2^{page_size+8} bytes; no single Serial Bus block read or block write transaction shall cross such a boundary. When *page_table_present* is zero, a *page_size* value of zero indicates that there are no alignment requirements. Nonzero *page_size* values specify alignment boundaries in the same fashion as when a page table is present.

The target shall issue data transfer requests with a speed equal to that specified by the *spd* field in the ORB. The target shall not issue block read or write requests with a data payload length greater than that specified by the *max_payload* field in the ORB.

Within the above speed, size and alignment constraints, the target is free to issue the data transfer requests in any order and to retry failed data transfer requests according to vendor-dependent algorithms.

9.3 Completion status

Upon completion of an ORB, the target shall examine the *notify* bit in the ORB to determine whether or not to store a status block. If *notify* is zero, the target may store a status block. Otherwise, if *notify* is one or if the ORB completed with an error condition, the target shall store a status block. The address for the status block is specified by *status_FIFO*, supplied by the initiator as part of the login or create stream request. The status block, previously described in 5.3, contains sufficient information to indicate successful command completion or, in the case of a faulted command, to permit the initiator to select the appropriate error handling strategy.

In all cases, the status FIFO allocated by the initiator shall be accessible to a single Serial Bus block write transaction with any *data_length* that is a multiple of four and less than or equal to 32 bytes. The target shall store the status block by means of a single block write and shall not attempt any retries if either:

- a) no acknowledge packet is received immediately after the write request; or

- b) subsequent to the receipt of an *ack_pending* immediately after the write request, no corresponding response packet is received within the split time-out limit.

Other errors, including the link layer busy conditions, *ack_data_error*, *resp_conflict_error* and *resp_data_error*, may be retried up to a vendor-dependent limit. If no retry is attempted or if the retry limit is exhausted without success, the target fetch agent shall transition to the DEAD state.

The return of completion status indicates to the initiator that the task commenced by the ORB is no longer part of the task set. The status block also specifies whether or not the system memory allocated to the ORB may be released. If the *src* field has a value of zero, the initiator may reuse or deallocate the system memory occupied by an ORB.

NOTE – For targets that support the ordered model of task execution, the return of completion status for an ORB implicitly indicates that all preceding ORB's in the linked list have completed successfully and are no longer part of the task set.

9.4 Unsolicited status

In addition to status associated with a particular ORB, described in the preceding section, a fetch agent may store unsolicited status at the address specified by *status_FIFO*. A status block that contains unsolicited status shall be identified by setting the *src* field to a value of two or three, according to whether the unsolicited information is device status or the report of an isochronous error, respectively (see 5.3).

A fetch agent may store unsolicited status at any time that its *unsolicited status enabled* variable is one. Upon successful completion of the Serial Bus block write transaction used to store the status block, the fetch agent shall zero its *unsolicited status enabled* variable. The initiator may set the fetch agent's *unsolicited status enabled* variable to one by writing any data value to the corresponding UNSOLICITED_STATUS_ENABLE register.

NOTE – One use for unsolicited status is to report progress of lengthy operations such as a disk format or tape rewind. Device implementers that use unsolicited status for this purpose should pick an appropriate interval for the reports. Frequent unsolicited status transfers reduce available Serial Bus bandwidth and may increase processing overhead at the initiator without any perceivable benefits.

The action taken by a target when unsolicited status is generated but cannot be stored because the *unsolicited status enable* variable is zero depends upon the nature of the status. If the status is for a unit attention condition, the target shall retain the information with the intent to store it as soon as the *unsolicited status enable* variable is set to one. The unit attention condition shall persist until the corresponding status block is stored at the initiator's *status_FIFO*. The definition of unit attention conditions is beyond the scope of SBP-2 and is usually the province of the command-set standard for the target. Other status information, that does not constitute a unit attention, may be discarded by the target, queued for future delivery or replace an existing, pending status of the same type. Which of these behaviors is appropriate is determined by the command set standard.

10 Task management

The preceding section describes the procedures used by the initiator to signal the target that tasks are to be executed and the procedures by which the target performs data transfer or device control for the tasks and ultimately signals their completion back to the initiator. Section 9 gives no consideration to the larger perspective of how these tasks interact with each other and how the initiator may manage the tasks.

This section defines how individual tasks are collected together as task sets and how both tasks and task sets may be managed by the initiator.

10.1 Task sets

A task set is a collection of tasks, each of which has an associated command in an ORB, that is available to the target for execution. The interactions among these tasks and the ordering relationships, if any, are governed by the task management model implemented by the target.

A task enters the task set when it is linked into an active request list. The extent of a task set includes all the uncompleted ORB's linked into a request list in system memory, not solely the ORB's already fetched by the target.

Historically, there has been one task set associated with each logical unit of a device. The concept of a task set is extended by SBP-2 to permit multiple stream task sets per logical unit. Each time target resources are allocated for isochronous operations (by means of a create stream request), a task set is created that is associated both with a logical unit and a stream identifier. There is a one-to-one relationship between a stream identifier and a stream task set; there may be multiple stream task sets associated with a logical unit. Each stream task set is separate and distinct from the normal task set and from other stream task sets: there are no interactions between tasks that belong to different stream task sets.

10.2 Basic task management model

Targets shall support, at a minimum, a basic task management model. Under this model, the following rules apply:

- All tasks within a task set share the same execution characteristics: either they are all reorderable or else they are all ordered;
- The reorderable or ordered execution characteristics of a task set are implicit in the target implementation and are not subject to control by the initiator;
- For stream task sets, the target shall execute all tasks in order and report their completion status in the same order. For normal task sets, configuration ROM shall specify whether the target may reorder task execution or not;
- All tasks within a task set are uniquely identified by the Serial Bus address of the ORB that initiated the task. This address shall be unique for the life of the task;
- The abort task, abort task set and target reset task management functions, described later in this section, shall be implemented;

The only element of choice in the implementation of a task set is under the basic model is whether or not the target may reorder task execution. An unordered model is usually appropriate for devices, such as mass storage, where no positional or other context information is inherited from one command to the next. An ordered model may be more appropriate for devices, such as sequential storage, where the outcome of one command affects the next. The same ordering considerations apply to stream task sets, within which the data is time-ordered by its very nature.

The unordered model is characterized by unrestricted reordering of the active tasks. The target may reorder the actual execution sequence of any tasks in a task set in any manner. Unrestricted reordering places the responsibility for the assurance of data integrity on the initiator. If the integrity of data on the device medium could be compromised by unrestricted reordering involving a set of active tasks, $\{T_0, T_1, T_2, \dots T_N\}$ and a new task T' , the initiator shall wait until $\{T_0, T_1, T_2, \dots T_N\}$ have completed before appending T' to an active request list.

NOTE – In multitasking operating system environments, independent execution threads may generate tasks that have ordering constraints within each thread but not with respect to other threads. If this is the case, an initiator may manage the constraints of each thread yet still keep the target substantially busy. This avoids the undesirable latencies that occur if the target is allowed to become idle before new ORB's are signaled.

The ordered model requires both that tasks are executed in order and that completion status is returned in order. Because Serial Bus transactions may complete out of order, the target shall single-thread the return of completion status. Once the target has transmitted a Serial Bus block write request to store completion status in system memory, it shall not attempt to store completion status for any other task in the task set until *ack_complete* or *resp_complete* is received.

10.3 Error conditions

Upon an error condition or fault detected during the execution of any task within a task set, the entire task set shall be cleared as follows:

- a) The target shall halt the operation of the fetch agent associated with the task set by making a transition to the DEAD state;
- b) For all recently completed tasks, the target shall wait until the completion status of each command has been successfully stored in system memory or until the implementation-dependent retry algorithms have been exhausted in the attempt to store completion status; and
- c) Finally, the target shall return error completion status for the faulted task. The *dead* bit shall be one in the status block.

The return of error status for a faulted task is an indication to the initiator that the task set has been cleared and that any remaining active tasks in the request list have been aborted.

10.4 Task management requests

The clauses that follow describe the use of the task management ORB's defined in 5.1.4.6.

10.4.1 Abort task

Abort task is a task management function that permits an initiator to abort a specified task. A modification to the *rq_fmt* field of the ORB to be aborted is the basic method; in addition, targets may also recognize task management ORB's to abort tasks. All targets shall support abort task.

Because the task to be aborted may not have been fetched by the target when the initiator wishes to abort the task, the following procedure shall be used to abort the task:

- a) The *rq_fmt* field shall be set to a value of three in the ORB for the task to be aborted. Note that this field and the *next_ORB* field are the only two portions of an ORB that may be modified by the initiator once the ORB is linked into an active request list;
- b) The initiator may construct a management ORB in system memory for the abort task function. The initiator shall set the appropriate values in the *rq_fmt*, *login_ID* and *ORB_offset* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to ABORT TASK; *ORB_offset* shall contain the Serial Bus address of the ORB for the task to be aborted;

- c) The initiator may signal the abort task management ORB to the management agent.

Mandatory support for abort task is dependent upon the target's ability to recognize an *rq_fmt* value of three in an ORB and take the actions described below.

- If the ORB to be aborted has already been fetched by the target, the task may be completed by the target without recognition of the abort task request; otherwise
- When the ORB is first fetched, the target shall recognize the *rq_fmt* field value of three and shall not execute the command. That target shall store completion status for the aborted ORB; the request status shall be REQUEST COMPLETE and the *sbp_status* field shall indicate dummy ORB completed.

A second method to abort task(s) may be available by means of task management ORB's with a *function* of ABORT TASK. Target support for this method of abort task is determined by the type of task set to which the task to be aborted belongs. The *login_ID* field in the ORB identifies both the fetch agent and the corresponding task set. If the *login_ID* was returned in a *login_response*, target support for this method of abort task is optional. Otherwise the *login_ID* was returned in a *create_stream_response* and the target shall implement support as specified below. Targets that support abort task by means of task management ORB's shall store a completion status of FUNCTION COMPLETE for the abort task request in the status buffer specified by the ORB.

If the task to be aborted, identified by *ORB_offset*, is not recognized by the target as part of its working set, one of two conditions may exist: either the ORB has not been fetched or completion status has already been stored. In either case the target is not required to take any immediate action. In the first case, when the ORB is ultimately fetched, the *rq_fmt* field has a value of three and the target shall not execute the command. The target shall store completion status for the aborted ORB; the request status shall be REQUEST COMPLETE and the *sbp_status* field shall indicate dummy ORB completed. In the second case, no action whatsoever need be taken by the target.

If the task to be aborted is recognized by the target as part of its working set, the target should attempt to abort the task according to the steps below. Note that timing conditions may exist that prevent targets from aborting the specified task. In particular, if the target has already issued a write request to store completion status for the task to be aborted, the target shall take no other action in response to the abort task request. Otherwise, the target should perform the following actions in response to a task management ORB with the ABORT TASK *function*:

- a) The target should not issue additional data transfer requests for the task;
- b) The target shall wait for responses to pending data transfer requests and, once all such responses are received, shall not issue additional data transfer requests for the task;
- c) So long as none of the target medium, data buffer or status FIFO have been modified as the result of partial execution of the task, the target shall store completion status of REQUEST COMPLETE with an *sbp_status* field that indicates dummy ORB completed;
- d) Otherwise, if task execution has commenced and any one of the target medium, data buffer or status FIFO has been modified, then the target shall store completion status of REQUEST with an *sbp_status* field that indicates request aborted.

Regardless of which abort task methods are supported by the target, the initiator shall not reuse the system memory occupied by the ORB, data buffer or page table of the task to be aborted until completion status is returned for that ORB.

10.4.2 Abort task set

Abort task set is a task management function that permits an initiator to abort all of its tasks within a task set. All targets shall support abort task set.

To abort a task set, the initiator shall construct a management ORB in system memory for the abort task set function. The initiator shall set the appropriate values in the *rq_fmt* and *login_ID* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to ABORT TASK SET.

The initiator shall signal the abort task set ORB to the management agent.

Upon receipt of an abort task set request, the target shall perform the following actions:

- a) The target shall halt the operation of the fetch agent associated with the task set by making a transition to the DEAD state;
- b) The target shall not issue data transfer requests for any task in the task set whose *login_ID* is equal to that specified in the abort task set request;
- c) The target shall wait for responses to pending data transfer requests for any task in the task set whose *login_ID* is equal to that specified in the abort task set request;
- d) For all recently completed tasks whose *login_ID* is equal to that specified in the abort task set request, the target shall wait until the completion status of each command has been successfully stored in system memory or until the implementation-dependent retry algorithms have been exhausted in the attempt to store completion status; and
- e) When all of the above events have completed, the target shall store completion status for the abort task set request in the status buffer provided. The completion status shall indicate FUNCTION COMPLETE.

The initiator shall not reuse the system memory occupied by any of the ORB's, data buffers or page tables of the tasks to be aborted until completion status is returned for the abort task set request.

10.4.3 Clear task set

Clear task set is a task management function that permits an initiator to abort all tasks within a task set. Targets may support clear task set.

To clear a task set, the initiator shall construct a management ORB in system memory for the clear task set function. The initiator shall set the appropriate values in the *rq_fmt* and *login_ID* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to CLEAR TASK SET.

The initiator shall signal the clear task set ORB to the management agent.

Upon receipt of a clear task set request, the target shall perform the following actions:

- a) The target shall halt the operation of all fetch agents associated with the task set by making transitions to the DEAD state;
- b) The target shall not issue data transfer requests for any task in the task set;
- c) The target shall wait for responses to pending data transfer requests for any task in the task set;
- d) For all recently completed tasks, the target shall wait until the completion status of each command has been successfully stored in system memory or until the implementation-dependent retry algorithms have been exhausted in the attempt to store completion status;
- e) The target shall create a unit attention condition for all logged-in initiators other than the initiator, identified by *login_ID*, that signaled the clear task set request; and
- f) When all of the above events have completed, the target shall store completion status for the clear task set request in the status buffer provided. The completion status shall indicate FUNCTION COMPLETE.

The initiator shall not reuse the system memory occupied by any of the ORB's, data buffers or page tables of the tasks to be aborted until completion status is returned for the clear task set request.

10.4.4 Logical unit reset

Logical unit reset is a task management function that causes a logical unit to perform the actions described below and to create unit attention conditions for all initiators logged-in to the logical unit. Support for logical unit reset is a target option.

To reset a logical unit, the initiator shall construct a management ORB in system memory for the logical unit reset function. The initiator shall set the appropriate values in the *rq_fmt* and *login_ID* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to LOGICAL UNIT RESET.

The initiator shall signal the logical unit reset ORB to the management agent.

Upon receipt of a logical unit reset request, the logical unit shall perform the following actions:

- a) The target shall halt the operation of all of the logical unit's fetch agents by making transitions to the DEAD state;
- b) The target shall not issue data transfer requests for any task in any of the logical unit's task sets;
- c) The target shall create a unit attention condition for all initiators logged-in to the logical unit other than the initiator, identified by *login_ID*, that signaled the logical unit reset request; and
- d) When all of the above events have completed, the target shall store completion status for the logical unit reset request in the status buffer provided. The completion status shall indicate FUNCTION COMPLETE.

The initiator shall not reuse the system memory occupied by any of the affected ORB's, data buffers or page tables of the tasks until completion status is returned for the target reset request.

10.4.5 Target reset

Target reset is a task management function that causes a target to perform the actions described below and to create unit attention conditions for all initiators. All targets shall support target reset.

To reset a target, the initiator shall construct a management ORB in system memory for the target reset function. The initiator shall set the appropriate values in the *rq_fmt* and *login_ID* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to TARGET RESET.

The initiator shall signal the target reset ORB to the management agent.

Upon receipt of a target reset request, the target shall perform the following actions:

- a) The target shall halt the operation of all fetch agents for all logical units by making transitions to the DEAD state;
- b) The target shall not issue data transfer requests for any task in any task set;
- c) The target shall create a unit attention condition for all logged-in initiators other than the initiator, identified by *login_ID*, that signaled the target reset request; and
- d) When all of the above events have completed, the target shall store completion status for the target reset request in the status buffer provided. The completion status shall indicate FUNCTION COMPLETE.

The initiator shall not reuse the system memory occupied by any of the affected ORB's, data buffers or page tables of the tasks until completion status is returned for the target reset request.

10.4.6 Terminate task

Terminate task is a task management function that permits an initiator to request early completion of a specified task. Targets that implement the basic task management model shall not support terminate task and shall reject all terminate task requests with a completion status of FUNCTION REJECTED.

To request task termination, the initiator shall construct a management ORB in system memory for the terminate task function. The initiator shall set the appropriate values in the *rq_fmt*, *login_ID* and *ORB_offset* fields of the ORB, as described in 5.1.4.6. The *function* field shall be set to TERMINATE TASK; *ORB_offset* shall contain the Serial Bus address of the ORB for the task to be terminated. Once the terminate task ORB has been initialized, the initiator shall signal the ORB to the management agent.

Upon receipt of a terminate task request, the target shall store a completion status of FUNCTION COMPLETE or FUNCTION REJECTED for the terminate task request in the status buffer provided.

If the terminate task function is accepted by the target, the completion status of FUNCTION COMPLETE does not necessarily indicate that the specified task has completed. The ultimate completion of the specified task shall be signaled when the target stores completion status for the task.

If an error condition is detected for the specified task, the terminate task request shall be ignored and the target shall perform the actions previously described in 10.3.

If the specified task completes prior to the receipt of the terminate task request, the target shall wait until completion status is successfully stored for the specified task before completion status shall be stored for the terminate task request.

Otherwise, the target shall complete the specified task as follows:

- a) The target shall not issue data transfer requests for the task;
- b) The target shall wait for responses to pending data transfer requests;
- c) The target shall store completion status of REQUEST COMPLETE and appropriate command set-dependent status that indicates command termination.

When a terminated task creates an error condition, the target shall clear the task set and take the actions described in 10.3.

The initiator shall not reuse the system memory occupied by the ORB, data buffer or page table of the task to be terminated until completion status is returned for that ORB.

10.5 Task management event matrix

Common events that affect the state of target fetch agents and their associated task set(s) are summarized below. Refer to the governing clauses in sections 8 and 9 as well as this section for detailed information.

Event	AGENT_STATE.st		Task set(s)	
	Normal	Stream	Normal	Stream
Power reset	RESET		Clear all task sets	
Command reset (write to RESET_START)	RESET		Clear all task sets	
Bus reset (immediate)	RESET	—	Clear all task sets	—
Bus reset (after one second)	—	—	Logout any initiator that has failed to successfully reconnect	
Login	—	—	—	—
Create stream	—	—	—	—
Reconnect	—	—	—	—
Logout	RESET		Abort initiator's task set	
Faulted command (CHECK CONDITION)	DEAD		Abort faulted initiator's task set	
ABORT TASK	—	—	—	—
ABORT TASK SET	DEAD		Abort initiator's task set	
CLEAR TASK SET	DEAD		Clear all task sets	
LOGICAL UNIT RESET	DEAD		Abort all the logical unit's task sets	
TARGET RESET	DEAD		Clear all task sets	
TERMINATE TASK	—	—	—	—

With respect to events supported by the target's management agent, *e.g.*, logout, there is an assumption of successful completion. In the case of a function rejected response or other indication of failure, the preceding table does not apply.

Bus resets affect target fetch agents and task sets according to the kind of request, login or create stream, by which the initiator first acquired access privileges. A login request allocates normal command block resources while a create stream request allocates stream command block and stream control resources.

Immediately upon detection of a bus reset, all normal command block fetch agents transition to the reset state and their associated task sets are cleared. Stream command block and stream control fetch agents do not fetch any additional ORB's subsequent to a bus reset but otherwise preserve their state. The task sets associated with these stream agents continue execution, but status for completed commands is held by the target and not stored to the initiator's *status_FIFO*.

For one second subsequent to a bus reset, targets save state information for initiators that were logged-in at the time of the bus reset. The one-second timer commences when the target observes the first subaction gap subsequent to a bus reset; if a bus reset occurs before the timer expires, the timer is reset. If an initiator successfully completes a reconnect request during this period, the actions described in 8.3 occur. For normal command block requests, the task set is empty and the initiator may signal new ORB's to the target. For both stream command block and stream control agents, fetching operations resume from the same point as before the bus reset. Any completion status held by the target during this one second period may also be stored to the initiator's *status_FIFO* after the successful reconnection.

One second after a bus reset, the target shall automatically perform a logout operation for all login ID's and stream ID's that have not been reconnected with their initiator. This returns all the affected fetch agents to the reset state and aborts any associated stream task sets.

11 Isochronous data interchange format

Isochronous data stored on the medium may be kept in a form similar to the format of isochronous packets on Serial Bus, but the *tcode* field present in Serial Bus packets is reused to identify the type of recorded isochronous data. Three different packet formats may be present in recorded isochronous data, encoded by *tcode* as shown below.

Value	Name	Description
8	CYCLE MARK	Marks the time of a cycle start event
A ₁₆	DATA	Isochronous data packet
E ₁₆	NULL	Null (or filler) packet
All other values	—	Reserved for future standardization

The values used to indicate CYCLE MARK and DATA are identical to the *tcode* values defined for Serial Bus cycle start packets and isochronous data packets, respectively.

In addition to the reuse of the *tcode*, the header and data CRC fields observed as part of Serial Bus isochronous packets are not recorded on the medium. Recorded isochronous data packets shall be stored on quadlet boundaries on the medium and shall contain an integral number of quadlets.

11.1 Cycle marks

Whenever a cycle start packet is observed on Serial Bus for an enabled isochronous stream, a CYCLE MARK packet shall be recorded on the medium. The CYCLE MARK packet is a single quadlet that stores the time transported by the cycle start packet, as shown by the figure below.




Figure 59 – CYCLE MARK format

The *second_count* and the *cycle_count* fields shall contain the values of the corresponding fields from the most recently observed cycle start packet. No more than one CYCLE MARK packet shall be recorded for a single cycle start packet.

NOTE – The time information in the CYCLE MARK packet is not necessary for a target to recreate an isochronous stream during playback, but it may be useful to applications that search for known time and cycle boundary locations in recorded isochronous data.

The *tcode* field shall be equal to eight.

When a target is a listener and detects a missed isochronous cycle, it shall synthesize and record a CYCLE MARK packet on the medium. The *second_count* and *cycle_count* values shall be taken from the target's free-running cycle timer.

11.2 Isochronous data packets

The format of an isochronous data packet recorded on the medium is illustrated below.

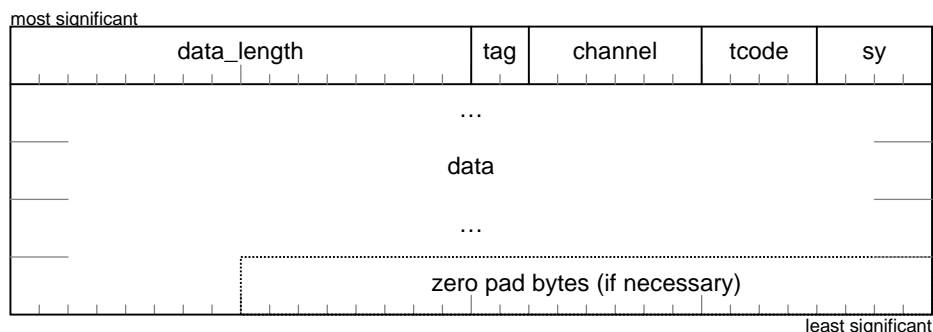


Figure 60 – Format for recorded isochronous data

The *data_length* field shall specify the length, in bytes, of the *data* field for the packet.

The *tag* field shall specify the format of the *data* field, encoded as indicated by the following table.

Value	Data field format
0	Unformatted data
1	Common isochronous packet (CIP) format
2 – 3	Reserved for future standardization

The *channel* field shall identify the isochronous channel number for the packet. The *channel* field recorded on the medium may have been transformed by the mapping from source channel (observed on Serial Bus) to *dest_channel* specified by a stream control ORB with a CONFIGURE CHANNELS control function (see 5.1.3). Upon playback, the *channel* field shall be transformed by the mapping to *dest_channel* specified by a stream control ORB with a CONFIGURE CHANNELS control function.

The *tcode* field shall be equal to A_{16} .

The *sy*, or synchronization code, field is an application-dependent field, the details of whose use are beyond the scope of this standard.

NOTE – A synchronization point may be defined as a boundary between video or audio frames, or any other point in the isochronous stream the application may consider appropriate.

The *data* field shall contain *data_length* bytes of information and shall be padded with trailing zero bytes, as necessary, to occupy an integral number of quadlets on the medium.

Dependent upon the value of *tag*, the target may require additional knowledge of isochronous data formats. When *tag* is zero the data payload of the isochronous packet is unformatted and requires no transformations upon either recording or playback. When *tag* is one, the format of the data payload shall conform to the common isochronous packet (CIP) format standardized by IEC 61883/FDIS.

11.3 Null packets

When the *tcode* field has a value of E_{16} , the data that is stored on the medium shall be ignored during playback. The format of a null packet is shown below.




Figure 61 – NULL packet format

The *data_length* field shall specify the length, in bytes, of the *data* field for the null packet. The number of quadlets occupied by the *data* field is $(data_length + 3)$ modulus 4. Zero is a permissible value for *data_length*; in this case, the null packet shall consist of only the header and shall be a single quadlet in length.

The *tcode* field shall be equal to E_{16} .

The values of quadlets within the *data* field are unspecified for NULL packets.

NOTE – NULL packets serve no particular purpose for targets, but they may be useful to some applications, such as nonlinear editing. Excessive quantities or sizes of NULL packets may cause some target implementations to experience underflow or other errors in the playback of isochronous data.

11.4 Cycle mark index

In addition to the data formats described above, a target may optionally record a cycle mark index as the last quadlet of every 512 bytes of isochronous data. A target shall not record a cycle mark index unless its medium has a block structure and the logical block size is a multiple of 512 bytes. The format of the cycle mark index is illustrated by Figure 62.




Figure 62 – Cycle mark index format

The *valid* bit (abbreviated as *v* in the figure above), when one, indicates that the *cm_index* field contains a valid cycle mark index.

The *cm_index* field, when valid, shall contain the position, in quadlets, of the first cycle mark in the 508 bytes of data that precede the cycle mark index. A *cm_index* value of zero specifies the first quadlet within the 512 byte region of data while a value of $7E_{16}$ specifies the quadlet that immediately precedes the cycle mark index.

NOTE – Targets may utilize *cm_index*, when valid, for error recovery during isochronous playback. An application that modifies recorded isochronous data shall insure the self-consistency of a 512-byte region and its cycle mark index when updated by normal command block ORB's.

11.5 Common isochronous packets (CIP)

Isochronous data packets (already described in 11.2) may also conform to a CIP format that divides the data payload into two parts: the CIP header and the application-dependent data that follows. Figure 63 illustrates the organization of the common isochronous packet format.




Figure 63 – Common isochronous packet (CIP) format

The CIP header is a variable number of quadlets (although only two are shown in the preceding figure). The most significant bit of each quadlet of the CIP header is called the *eah* bit. For an n quadlet CIP header, *eah* shall be zero for quadlets zero through $n - 2$, inclusive, and *eah* shall be one for quadlet $n - 1$. The next most significant bit of each quadlet is called the *form* bit. Together, the *eah* and *form* bits specify the format of the CIP header quadlet. CIP header formats are defined for *form* values of zero; *form* values of one are reserved for future standardization.

The only CIP header format currently defined is a two-quadlet header shown below.




Figure 64 – Two-quadlet CIP header format

The *sid*, or source ID, field shall specify the Serial Bus physical ID of the source (talker) for the isochronous data. Upon playback, the target shall substitute the *sid* value specified by the channel configuration map (see 5.1.3) for the recorded *sid* value in all transmitted isochronous packets.

The *dbs*, or data block size, field shall specify the size of each of the application-dependent data block(s) that follow the CIP header. A *dbs* value of zero encodes a size of 256 quadlets; for all other values of *dbs* the number of quadlets is the value of *dbs* itself. More than one data block may be encapsulated within a single Serial Bus isochronous packet.

The *fn*, or fraction number, field shall specify the number of data blocks that form a higher level, application-dependent object—the isochronous source packet. The number of data blocks that form an isochronous source packet is specified as 2^{fn} ; when there is a one-to-one correspondence between isochronous source packets and data blocks *fn* shall have a value of zero.

The *qpc*, or quadlet padding count, field shall specify the number of pad quadlets appended to an isochronous source packet before it is divided into data blocks. The quadlet padding count shall be less than the data block size and shall specify a value that results in equal sizes for the data blocks. If *fn* is zero, *qpc* shall also be zero. When *qpc* is nonzero the last data block shall include the pad quadlets, which are transmitted when the target is a talker.

The *sph*, or source packet header, bit (abbreviated as *s* in Figure 64) shall be one if the isochronous source packet begins with a header quadlet of the format shown below; otherwise, it shall be zero.




Figure 65 – Source packet header format

The source packet header shall contain a time stamp encoded in the same fashion as the least significant 25 bits of the CYCLE TIME register.

The *dbc*, or data block continuity counter, field shall specify the sequence number of the isochronous source packet and the sequence number of the data block within the isochronous source packet. The least significant *fn* bits of *dbc* hold the sequence number of the data block while the most significant $8 - fn$ bits hold the sequence number of the isochronous source packet itself. The data block continuity counter labels the first data block that follows the CIP header; the continuity counter of additional data blocks after the first increases monotonically from the value of *dbc*.

NOTE – The data block that immediately follows the CIP header is not necessarily the first data block of the isochronous source packet. The location of the starting data block of an isochronous source packet can be determined from the values of *dbc* and *fn*. Relative to the first data block after the CIP header (counting from zero), the ordinal of this data block is given by $(2^{fn} - (dbc \bmod 2^{fn})) \bmod 2^{fn}$. If a source packet header is present (as indicated by the *sph* bit), it shall be the first quadlet of this data block.

The *fmt* field shall specify the formats of both the *fmt*-dependent field within the same quadlet of the CIP header and the application-dependent data contained within the common isochronous packets. An *fmt* value of $3F_{16}$ indicates that no application-dependent data follows the CIP header and that the *dbc*, *fn*, *qpc* fields, the *sph* bit and the *dbc* field in the CIP header shall all be ignored. Other values of *fmt* encode the application-dependent format of the isochronous data, *e.g.*, DVCR or MPEG. The details of most application-dependent formats are not relevant to targets and are beyond the scope of this standard. However, the most significant bit of *fmt* specifies the format of the *fmt*-dependent field within same quadlet of the CIP header; this field is meaningful to targets when it contains a time stamp, since the time stamp shall be transformed during playback. The table below summarizes the meanings of *fmt* for targets.

Value	Description
0 – $1F_{16}$	Application data is present; the <i>fmt</i> -dependent field contains a time stamp defined by <i>syt</i> below
20_{16} – $3E_{16}$	Application data is present; the contents of the <i>fmt</i> -dependent field are unspecified
$3F_{16}$	No application data is present

When *fmt* is in the range zero to $1F_{16}$, inclusive, the second quadlet of the CIP header has the format illustrated below.




Figure 66 – Synchronization time (*syt*) format

The two fields, *cycle_count* and *cycle_offset*, are collectively referred to as the *syt*, or synchronization time, field. When *syt* has a value of FFFF_{16} , no synchronization time information is present and the *syt* field value shall be preserved upon playback. Otherwise, the *syt* field represents a time stamp encoded in the same fashion as the least significant 16 bits of the CYCLE_TIME register. Just as in the case of the CYCLE_TIME register, the value of *cycle_offset* is constrained to be in the range zero to 3071 inclusive; Values of *syt* for which *cycle_offset* is greater than 3071 are invalid. When *syt* contains valid *cycle_count* and *cycle_offset* fields, the target shall transform these values upon playback, as described in 12.2.3.

12 Isochronous operations

For each active channel on Serial Bus, isochronous data consists of zero or one isochronous packet transmitted by a talker in an isochronous cycle and received by zero or more listeners in the same isochronous cycle. This section describes how an initiator may control isochronous data transfers when a target participates, as either the talker or a listener, with other isochronous device(s) on Serial Bus.

Control of isochronous streams involves the following elements:

- the allocation of target resources (create stream requests);
- the establishment or breaking of connections between the target the target and other isochronous devices (connection management);
- the transfer of isochronous data to or from the target's medium (stream command block requests);
- the starting, stopping and synchronization of isochronous data reception or transmission by the target from or to Serial Bus (stream control ORB's); and
- the allocation of Serial Bus resources, such as channel numbers or bandwidth.

The first item, create stream, has already been defined as part of the access procedures in section 8. The second, connection management, may not affect target implementations and is outside of the scope of this document. The last, channel and bandwidth allocation, is specified by IEEE Std 1394-1995. This section describes the remaining procedures of isochronous operations.

12.1 Stream command block requests

With the notable exception of the lack of *data_descriptor* and *data_size* fields, stream command block requests are similar to and operate in a like fashion as normal command block requests that transfer data to or from the medium.

The target agent that is responsible to fetch stream command block ORB's operates in the same way as the command block agent(s). The initiator may build linked lists of stream command block ORB's, signal them to the target and dynamically append new stream command block ORB's while the target is active. See section 9 for a more detailed description; the information is equally applicable to normal command block requests and stream command block ORB's.

Stream command block requests shall be completed in order. The target shall insure that data is transferred in the same order as specified by the linked list of stream command block ORB's. Each isochronous stream active at a logical unit shall have its own task set that is disjoint from all other task sets for the logical unit. The task set for an isochronous stream is instantiated when a successful create stream request completes. The details of a task set for isochronous streams are explained at length in 10.1.

A stream command block ORB that transfers data from the device medium shall return completion status to the initiator when *stream_length* bytes of data have been successfully transferred or an error condition occurs. The return of completion status to the initiator does not provide any information as to how many of the *stream_length* bytes of data have been transmitted on Serial Bus.

A stream command block ORB that transfers data to the device medium shall return completion status to the initiator when *stream_length* bytes of data have been successfully transferred or an error condition occurs.

NOTE – A bus reset can prevent the target from fetching new stream command block requests from the initiator for up to one second—until the initiator successfully completes a reconnection. In order to provide for the uninterrupted flow of isochronous data when a bus reset occurs, the initiator should insure that sum of

stream_length for all the uncompleted stream command block ORB's represents a minimum of slightly more than one second's isochronous data at the configured data rate(s).

12.2 Stream control

The stream controller is implemented within the target to mediate the transfer of data between Serial Bus and the medium; the stream command block ORB's contain instructions that govern this transfer. The stream controller's responsibilities include:

- the synchronization, starting and stopping of either the reception or transmission of isochronous data from or to Serial Bus. The synchronization may occur at a specified time or may occur in response to a specified data pattern in observed isochronous data;
- the selective enablement of specified isochronous channels. When listening, the stream controller selects which channels to receive and provide to the stream command block ORB's. When talking, the stream controller selects which channels to transmit from the data provided by stream command block ORB's; and
- the transformation of Serial Bus header and common isochronous packet (CIP) header information between the representation of the isochronous data on Serial Bus and on the device medium. This may include the channel number and time-stamp information and the generation or elimination of null packets, as appropriate.

The target agent that is responsible to fetch stream control ORB's supports a linked list of requests that the initiator may build in system memory. This capability is essential for stream control ORB's, since synchronization boundaries may occur with as little separation as 125 μ s. See section 9 for a more detailed description; the information is as applicable to stream control ORB's as it is to normal command block requests and stream command block ORB's.

The clauses below describe the procedures that an initiator shall use to govern the actions of the stream controller. In addition, the stream controller shall transform isochronous header information independent of any stream control ORB's issued by the initiator.

12.2.1 Channel masks

Targets are configured to receive or transmit isochronous channels independently of the selective enablement or disablement of the individual channels. The selection of an active set of channels shall be performed by the stream controller according to a channel mask maintained for each stream.

The channel mask may be updated by a stream control ORB with a SET CHANNEL MASK control function, as described in 5.1.3. The channel mask may be updated only when the stream controller is stopped or paused. At all other times the target shall reject any attempts to modify the channel mask and shall leave it unchanged.

When the target is a listener, the channel mask refers to channel numbers as they are observed on Serial Bus. The channel mask is applied to select isochronous packets from Serial Bus before any transformations are performed and before the data is recorded on the medium.

When the target is a talker, the channel mask refers to channel numbers as they are recorded on the medium. The channel mask is applied to the stream of isochronous data presented to the stream controller before any transformations are performed and before isochronous packets are transmitted on Serial Bus.

12.2.2 Flow control

The execution of the stream control ORB functions START, STOP, PAUSE and UPDATE CHANNEL MASK can be synchronized with external events, as described in 5.1.3. This section describes the effect of synchronization on the transfer of data between the stream controller and Serial Bus.

When the target is a talker, data flows from the device medium (under the control of stream command block ORB's) to the stream controller. If the stream controller is in a state where isochronous data transmission is enabled (typically as the result of a stream control ORB with a START control function), isochronous data is selected according to the stream's channel mask, is then transformed as described in 12.2.3 and transmitted on Serial Bus. If the stream controller is in a state where isochronous data transmissions are disabled (typically as the result of a stream control ORB with a STOP or PAUSE control function), isochronous data transfer from the medium may continue within the limitations of target buffer resources. Once these buffers are full, the execution of stream command block ORB's shall be suspended.

When the target is a listener, data flows from Serial Bus to the stream controller. If the stream controller is in a state where isochronous data reception is enabled (typically as the result of a stream control ORB with a START control function), isochronous data is selected according to the stream's channel mask, is then transformed as described in 12.2.3 and is transferred to the device medium under the control of stream command block ORB's.

12.2.3 Isochronous data transformation

The stream controller is responsible to both filter and transform isochronous data as it is recorded on or played from the device medium. Filtering is performed according to the channel number(s) of the isochronous packet(s) (either received from Serial Bus or read from the device medium) before any transformations are applied. Once a set of enabled channel(s) is selected, the stream controller shall:

- replace the channel number found in the source isochronous data with the remapped channel number for the destination isochronous data. This may be the identity map;
- upon playback, if the isochronous data conforms to the common isochronous packet (CIP) format, replace the *sid* (or source ID) field in the CIP header with the source ID specified by a CONFIGURE CHANNELS request;
- when recording, transform all instances of time stamps in the CIP or source packet headers to reflect the delta, or time difference, between the time stamp and the Serial Bus isochronous cycle time at which the data is observed; and
- upon playback, add the Serial Bus isochronous cycle time at which the data is transmitted to all instances of time stamps in the CIP or source packet headers, in order to recreate an absolute time stamp.

The first two transformations are straightforward and require no additional explanation. The transformations to be applied to time stamps are more complex and are explained below.

The common isochronous packet (CIP) format, as standardized by IEC 61883/FDIS, may contain isochronous time stamp information that is absolute rather than relative. That is, the time stamps contained within header data are a fixed offset ahead of, in the most significant bits, the isochronous cycle times contained in the cycle start packet that indicates the cycle in which the packets are transmitted. Time stamps are found in two places in packets that conform to CIP format:

- the *syt* field of the second quadlet of a two-quadlet CIP header if the *fmt* field in that quadlet has a value between zero and 1F₁₆, inclusive; and
- the *cycle_count* field of the isochronous source packet header.

See 11.3 for the exact specifications of these fields and the circumstances under which they are present in isochronous data.

In order to permit the subsequent recreation of absolute time stamps upon playback, the stream controller shall calculate and store the delta, or time difference, between the time stamp in the observed data and the CYCLE_TIME register when the data is observed.

For the *syt* field, the delta time shall be obtained by applying the following formula (shown in C code notation):

$$syt_{\text{stored}} = syt_{\text{observed}} - (\text{CYCLE_TIME} \& 0x0000F000);$$

For the *cycle_count* component of source packet headers, the delta time shall be obtained by applying the following formula:

$$\text{cycle_count}_{\text{stored}} = (8000 + \text{cycle_count}_{\text{observed}} - \text{CYCLE_TIME.cycle_count}) \% 8000;$$

To transform the *syt* field in a CIP header upon playback, the value of the CYCLE_TIME register for the isochronous cycle in which the packet is transmitted shall be masked with 01FF F000₁₆ and the result shall be added to the time stamp obtained from the medium. The most significant bits of the sum shall be discarded so the result fits within the 16-bit *syt* field. The resultant value, now an absolute time in terms of current isochronous cycle time, shall be transmitted in the outbound isochronous packet.

Because IEEE Std 1394-1995 constrains *cycle_count* to the range zero to 7999, inclusive, the playback transformation of the *cycle_count* component of the source packet header differs. The addition of a delta time to the *cycle_count* shall be performed modulus 8000, as shown below:

$$\text{cycle_count}_{\text{transmitted}} = (\text{cycle_count}_{\text{stored}} + \text{CYCLE_TIME.cycle_count}) \% 8000;$$

12.3 Error logs

When an isochronous stream is active, a stream controller may detect error conditions on Serial Bus or internally. Typical errors include but are not limited to:

- a missing isochronous packet or cycle start indication;
- an isochronous packet with a data CRC error;
- when the target is a talker, an underflow in the availability of data from the stream command block ORB's that causes no data to be transmitted for one or more channels during an isochronous cycle; or
- when the target is a listener, an overflow in which isochronous data from Serial Bus must be discarded because of an internal buffer overflow or a lack of stream command block ORB(s) to transfer the data to the medium.

The stream controller reports these and other errors by storing unsolicited status at the *status_FIFO* address specified at the time the stream was created. The stream controller may be placed in one of three error reporting modes by means of a stream control ORB with a SET ERROR MODE control function:

- report the first error and stop execution of the current stream control ORB;
- report all errors but continue executing the current stream control ORB; or
- ignore all errors and continue executing the current stream control ORB.

The *rpt* field in the stream control ORB establishes one of the three error reporting modes. See 5.1.3 for details of the SET ERROR MODE control function.

Upon detection of an isochronous error while talking or listening, if error reporting has been enabled the target shall store the status block shown below at the *status_FIFO* address provided by the initiator as part of the create stream request (see 5.1.4.3).




Figure 67 – Unsolicited status format for isochronous errors

The target shall store eight bytes of status information and shall use a single Serial Bus block write transaction to store the status block at the *status FIFO* address.

The *src* field shall be three to indicate that the status block describes an isochronous error.

The *resp* field shall be one to indicate a response status of TRANSPORT FAILURE.

The *len* field shall specify the quantity of valid status block information stored at the *status_FIFO* address and shall be equal to one.

The *stream_error* field shall specify the nature of the isochronous error, as encoded by the table below.

Value	Stream error description
0	Reserved (not to be used)
1	Missing CYCLE START packet
2	Data CRC error in received isochronous packet
3	Data length error in received isochronous packet
4	Internal underflow with the result that recorded isochronous data was not transmitted on Serial Bus
5	Internal overflow with the result that isochronous data observed on Serial Bus was not recorded on the medium
6 — FE ₁₆	Reserved for future standardization
FF ₁₆	Unspecified error

The *seconds* field shall contain the least significant 19 bits of the BUS_TIME register at the time of the isochronous stream error.

The *cycle_count* field shall contain the cycle count, between zero and 7999, at the time of the error. The cycle count shall be obtained from the target's free-running cycle timer and shall not be latched from the last observed CYCLE START packet.

Annex A (normative)

Minimum Serial Bus node capabilities

In addition to those minimum capabilities defined by IEEE Std 1394-1995, this annex specifies the minimum capabilities that an initiator or a target shall support in order to implement SBP-2.

A.1 Initiator capabilities

An initiator shall be capable of responding to block read or write requests with a *data_length* less than or equal to 32 bytes.

An initiator shall also be capable of responding to block read requests with a *data_length* less than or equal to $4 * ORB_size$, where *ORB_size* is obtained from the Logical_Unit_Characteristics entry in the target's configuration ROM.

For the largest value of *max_payload* specified in any normal command block ORB signaled to the target, the initiator shall be capable of responding to block read and write requests with a *data_length* less than or equal to $2^{max_payload + 2}$ bytes.

The initiator shall report the largest of these possible *data_length* values by setting the value of the *max_rec* field in the bus information block in its configuration ROM to a value of $(\log_2 data_length) - 1$.

A.2 Target capabilities

A target shall be capable of responding to block read or write requests with a *data_length* equal to eight bytes if the *destination_offset* specifies either the MANAGEMENT_AGENT or the ORB_POINTER register.

A target shall be capable of initiating block write requests with a *data_length* of at least eight bytes. Consequently, the *dreq* bit in the STATE_CLEAR and STATE_SET registers shall be implemented.

The target shall report this capability by setting the *drq* bit in the Node_Capabilities entry in configuration ROM to one.

NOTE – The value of STATE_CLEAR.*dreq* shall be unaffected by a Serial Bus reset. The target may automatically set *dreq* to zero (request initiation enabled) upon a power reset or a command reset.

The target shall report this capability by setting the value of the *max_rec* field in the bus information block in configuration ROM to a value of two.

Once a target completes its power reset initialization sequence, it shall acknowledge, and subsequently respond to, Serial Bus transaction request subactions within the time limits specified by IEEE Std 1394-1995 and draft standard P1394a. A Serial Bus reset, by itself, shall not alter a target's responsiveness to request subactions.

NOTE – While initializing after a power reset, a target shall respond to quadlet read requests addressed to FFFF F000 0400₁₆ with either a response data value of zero or acknowledge the request subaction with *ack_tardy*, as specified by draft standard IEEE P1394a. This indicates that the remainder of configuration ROM, as well as other target CSR's, are not accessible.

A.3 Target security

As mandated by draft standard P1394a, a target shall abide by the following restrictions:

- If a target's unique ID, EUI-64, is read from the configuration ROM bus information block by quadlet read requests, the value returned shall be the EUI-64 assigned by the manufacturer;
- The target shall not originate an asynchronous request or response packet with a *source_ID* field that is not equal to either a) the most significant 16 bits of the target's NODE_IDS register or b) the concatenation of $3FF_{16}$ and the physical ID assigned to the target's PHY during the self-identify process; and
- The target shall not receive an asynchronous request or response packet unless the *destination_ID* field is equal to either a) the concatenation of the most significant 10 bits of the target's NODE_IDS register and either the physical ID assigned to the target's PHY during the self-identify process or $3F_{16}$, or b) the concatenation of $3FF_{16}$ and either the physical ID assigned to the target's PHY during the self-identify process or $3F_{16}$.

Annex B (normative)

SCSI command and status encapsulation

SBP-2 defines a protocol that permits initiator(s) to control the operation of devices (disks, tapes, printers, *etc.*), but it does not specify the command sets used by the devices—only the mechanisms by which commands, data and status are transported. This annex specifies how SBP-2 may be used for devices that use the SCSI command set(s). This encompasses encapsulation of SCSI command descriptor blocks (CDB's), a standard format for SCSI status and sense data and the necessary configuration ROM entries.

B.1 SCSI command descriptor block

SBP-2 provides for the transport of 6-, 10- and 12-byte SCSI CDB's within a normal command block ORB, as illustrated by Figure B-1. When 6- and 10-byte CDB's are encapsulated within an ORB, the least significant (unused) bytes of the ORB shall be zero.




Figure B-1 – SCSI command block ORB

A SCSI CDB may also be transported within a stream command block ORB (see 5.1.2.2).

Parts of the control byte (the last byte of a SCSI command descriptor block) are constrained to values illustrated by Figure B-2.




Figure B-2 – SCSI control byte

The *naca*, or normal ACA, bit shall be zero; SBP-2 supports SCSI-2 contingent allegiance.

The *flag* and *link* bits shall be zero; SBP-2 does not support linked commands.

B.2 SCSI status and sense data

Upon completion of a command, if the *notify* bit in the ORB is one or if there is exception status to report, the target shall signal the initiator by storing all or part of the status block shown by Figure B-3 at the *status_FIFO* address provided by the initiator as part of the login request.




Figure B-3 – Status block format for SCSI sense data

When a command completes with GOOD status, only the first two quadlets of the status block shall be stored at the *status_FIFO* address; the *len* field shall be one. Otherwise, both SCSI status and sense data shall be stored in a status block that conforms to the format illustrated above.

NOTE – SBP-2 permits the return of a status block between two and eight quadlets in length. When a truncated status block is stored, the omitted quadlets shall be interpreted as if zero values were stored.

The *src*, *resp*, *len*, *sbp_status*, *ORB_offset_hi* and *ORB_offset_lo* fields, as well as the *dead* bit (abbreviated as *d* in the figure above), are as previously described in 5.3.

The *sfmt* field shall specify the format of the status block and shall additionally indicate whether the error condition associated with *sense_key* is current or deferred. The table below defines permissible values for *sfmt*.

Value	Description
0	Current error; status block format defined by this standard
1	Deferred error; status block format defined by this standard
2	Reserved for future standardization
3	Status block format vendor-dependent

The *status* field shall specify SCSI status information as defined by SAM-2, with the exceptions noted in the table below.

Value	Description
0	GOOD
2	CHECK CONDITION
4	CONDITION MET
8	BUSY
10 ₁₆	Not supported by SBP-2
14 ₁₆	Not supported by SBP-2
18 ₁₆	RESERVATION CONFLICT
22 ₁₆	COMMAND TERMINATED
28 ₁₆	Not supported by SBP-2
30 ₁₆	Not supported by SBP-2
All other values	Reserved for future standardization

The *valid* bit (abbreviated as *v* in the figure above) shall specify the content of the *information* field. When the *valid* bit is zero, the contents of the information field are not specified. When the *sfmt* field has a value of zero or one and the *valid* bit is one, the contents of the information field shall be as defined by SPC-2 or the relevant command set standard.

The meanings of the *mark*, *eom* and *illegal_length_indicator* bits (abbreviated as *m*, *e* and *i*, respectively, in the figure above) are defined by SPC-2 or the relevant the command set standard. These bits correspond to the filemark, EOM and ILI bits defined by SPC-2 for sense data.

The *sense_key*, *sense_code* and *sense_qualifier* fields shall specify command completion information defined by SPC-2 or the relevant the command set standard. These fields correspond to the sense key, additional sense code and additional sense code qualifier fields defined by SPC-2 for sense data.

The contents of the *information* field are unspecified if either the *valid* bit is zero or the *sfmt* field has a value of three. For *sfmt* values of one or two, the contents of the *information* field are device-type or command dependent and, if the *valid* bit is one, are defined within SPC-2 or the appropriate standard for the command. Characteristic uses of the *information* field are for:

- the unsigned logical block address associated with *sense_key* and the command; or
- the least significant 32-bits of the unsigned logical block address associated with *sense_key* and the command; or
- the residue of the requested data transfer length minus the actual data transfer length, in either bytes or blocks as determined by the command. Negative values are indicated in two's complement notation.

The contents of the *CDB*-dependent field are device-type or command dependent and are defined within SPC-2 or the appropriate standard for the command.

Nonzero values in the *fru* field may be used to identify a device-dependent, field replaceable mechanism or unit that has failed. A value of zero in this field shall indicate that no specific mechanism or unit has been identified to have failed or that the data is unavailable. When *fru* is nonzero, the format of the information is not specified by this standard.

When *sfmt* has a value of zero or one, the contents of the *sense_key*-dependent field are defined by SPC-2 or the relevant the command set standard. In this case the most significant bit of the *sense_key*-dependent field is the SKSV bit defined by SPC-2. When *sfmt* is equal to three, the contents of *sense_key*-dependent are unspecified.

B.3 Configuration ROM

SCSI targets shall implement configuration ROM in accordance with section 7 and this annex. At least one logical unit, logical unit zero, shall be implemented; additional logical units may be implemented. A logical unit is described by entries in a unit directory or by entries in a logical unit directory dependent upon the unit directory or by entries taken in combination from both places.

Mandatory and optional configuration ROM entries for SCSI targets, and the directories in which they may occur, are summarized by the table below.

Entry	Unit directory	Logical unit directory	Comments
Unit_Spec_ID	Mandatory	—	
Unit_SW_Version	Mandatory	—	
Command_Set_Spec_ID	Mandatory (in at least one directory)		See precedence rules below.
Command_Set	Mandatory (in at least one directory)		See precedence rules below.
Command_Set_Revision	Optional	Optional	See precedence rules below.
Firmware_Revision	Optional	—	
Management_Agent	Mandatory	—	
Logical_Unit_Characteristics	Mandatory (in at least one directory)		See precedence rules below
Logical_Unit_Directory	Optional	—	
Logical_Unit_Number	Mandatory (in at least one directory)		See precedence rules below.
Unit_Unique_ID	Optional	—	

For the entries that may be present in either a unit directory or a logical unit directory, the entry in the logical unit directory takes precedence over that in the parent unit directory (if any). If there is no entry in the logical unit directory, the value is inherited from the corresponding entry in the parent unit directory.

The presence of one or more Logical_Unit_Number entries in a unit directory makes the Command_Set_Spec_ID, Command_Set and Logical_Unit_Characteristics entries mandatory in the unit directory.

Within a logical unit directory, at least one Logical_Unit_Number entry is required and the Command_Set_Spec_ID, Command_Set and Logical_Unit_Characteristics entries are all mandatory unless the corresponding entry is present in the parent unit directory.

B.3.1 Command_Set_Spec_ID entry

The Command_Set_Spec_ID entry is an immediate entry in either a unit or logical unit directory that specifies the organization responsible for the command set definition for the target. The format of this entry is specified by 7.4.3.

SCSI targets shall have a *command_set_spec_ID* value of 00 609E₁₆, which indicates that NCITS is responsible for the command set definition.

B.3.2 Command_Set entry

The Command_Set entry is an immediate entry in either a unit or logical unit directory that, in combination with the *command_set_spec_ID*, specifies the command set implemented by the target. The format of this entry is specified by 7.4.4.

SCSI targets shall have a *command_set* value of 01 04D8₁₆, which indicates that the target's command set is specified by SCSI Primary Commands 2 (SPC-2) and related command set standard(s)—as determined by the target's peripheral device type(s). In addition, this *command_set* value specifies that the target conforms to all requirements of this annex.

B.3.3 Logical_Unit_Number entry

The Logical_Unit_Number entry is an immediate entry in either a unit or logical unit directory that specifies the peripheral device type and logical unit number of a logical unit implemented by the SCSI target. The format of this entry is specified by 7.4.10.

The *device_type* field indicates the peripheral device type implemented by the logical unit. This field shall contain a value specified by the table below.

Value	Peripheral device type
0 – 1E ₁₆	The value of <i>device_type</i> shall have the same meaning as the peripheral device type field returned in INQUIRY data as specified by SPC-2
1F ₁₆	Unknown device type

Annex C (normative)

Security extensions

SBP-2 specifies an access protocol, in section 8, that by itself makes no provisions for security. This annex defines extensions to SBP-2 that may be implemented by targets to provide some measure of security. Targets that implement these security extensions shall conform to all provisions of this annex.

Conformance to this annex does not preclude additional, command set-dependent security facilities.

C.1 Passwords

A target shall implement two passwords:

- The master password, which shall be unchangeable and equal to the target serial number. The target serial number should be in a humanly readable form affixed to the target. The master password shall not be readable *via* the target's Serial Bus interface except by a logged-in initiator; and
- The current password, which shall accommodate 28 bytes of password data and shall be alterable only by the set password function (see C.3).

All password values shall be unchanged by power reset, bus reset or command reset.

The value of the master password shall be obtainable by command set-dependent means.

A target may be manufactured with a current password of all zeros, with the expectation that the user assign a nonzero current password as part of target initialization. If a target is manufactured with a nonzero current password, the target shall be shipped with the current password in a humanly readable form.

C.2 Login

The description of the login protocol below reproduces that specified by section 8 and adds validation of cumulative login attempts and the *password* field from the login request. The target shall implement an internal counter, *login_attempts*, which shall be zeroed upon a power reset or upon a successful login or logout request. The target shall perform the following to process a login request:

- a) The target shall reject the login request if *login_attempts* is equal to three;
- b) The target shall read the initiator's unique ID, EUI-64, from the bus information block by means of two quadlet read transactions. The *source_ID* from the write transaction used to signal the login ORB to the target's MANAGEMENT_AGENT register shall be used as the *destination_ID* in the quadlet read transactions;
- c) The target shall determine whether or not the initiator already owns a login by comparing the EUI-64 just obtained against the *login_owner_EUI_64* for all *login_descriptors*. If the initiator is currently logged-in to the same logical unit, the login request shall be rejected but *login_attempts* shall not be incremented;
- d) The target shall validate the password provided by the login request. If *password_length* is zero, the password is eight bytes of immediate data present in the *password* field. Otherwise *password_length* specifies the size of the password addressed by *password*. If *password_length* is greater than 28 the target shall increment *login_attempts* and reject the login request. When *password_length* is valid,

the password provided is extended to 28 bytes by the addition of least significant bytes of zeros; the result is compared with the target's passwords. If the password provided fails to match either the target's current or master password, the *login_attempts* shall be incremented and the login request shall be rejected;

- e) If the *exclusive* bit is set in the login ORB, the target shall reject the login request if there are any active *login_descriptors* for the logical unit but shall not increment *login_attempts*;
- f) If an active *login_descriptor* with the *exclusive* attribute exists for the *lun* specified in the login ORB, the target shall reject the login request but shall not increment *login_attempts*; else
- g) The target shall determine if a free *login_descriptor* is available. If a *login_descriptor* is free, the initiator's *source_ID* is stored in *login_owner_ID*, the initiator's EUI-64 is stored in *login_owner_EUI_64*, the *lun* and *status_FIFO* fields from the login ORB are stored in the *login_descriptor*, the *exclusive* variable in the *login_descriptor* is set to the value of the *exclusive* bit from the login ORB and the addresses of the fetch agent(s) are stored in the *login_descriptor*. Lastly the target assigns a unique *login_ID* to this login and stores it in the *login_descriptor*.

If the target is able to satisfy the login request, it shall zero *login_attempts* and return a login response as specified in 5.1.4.1.

C.3 Set password

In order to change a target's current password, an initiator may use a management ORB with the format shown below.




Figure C-1 – Set password ORB

The *password* and *password_length* fields specify the new value for the current password. If *password_length* is zero, the *password* field contains immediate data. When *password_length* is nonzero, the *password* field shall conform to the format for address pointers specified by Figure 11 and shall specify the address of a buffer. The maximum value of *password_length* shall be 28. The buffer shall be in the same node as the initiator and shall be accessible to a Serial Bus block read request with a data transfer length less than or equal to *password_length*.

The *notify* bit and the *rq_fmt* and *function* fields are as previously defined for management ORB formats.

The *login_ID* field shall specify a login ID value obtained as the result of a successful login.

The *status_FIFO* field shall specify an address allocated for the return of status for the SET PASSWORD request, only. The contents of this field shall not update the status FIFO address established by the successful login that returned *login_ID*.

If *login_ID* specifies a valid current login for the initiator that signaled the SET PASSWORD request to the target's MANAGEMENT_AGENT register, the target shall update the current password to the new value specified by the set password request. The target shall not return completion status for the request unless either the request is rejected or the new password has been successfully stored such that it will not be affected by any subsequent power reset, bus reset or command reset.

Annex D (informative)

Sample configuration ROM

Configuration ROM is located at a base address of FFFF F000 0400₁₆ within a node's initial memory space. The requirements for general format configuration ROM for targets are specified in section 7. This annex contains illustrations of typical configuration ROM for a simple targets.

D.1 Basic target

Figure D-1 below shows the bus information block, root directory and unit directory for a basic SBP-2 target that implements a single logical unit.

most significant		
4	11 ₁₆	ROM CRC (calculated)
3133 3934 ₁₆ (ASCII "1394")		
node_options (00FF 2000 ₁₆)		
node_vendor_ID		chip_ID_hi
chip_ID_lo		
4	Root directory CRC (calculated)	
03 ₁₆	module_vendor_ID	
81 ₁₆	Text leaf offset	
0C ₁₆	node_capabilities (00 83C0 ₁₆)	
D1 ₁₆	Unit directory offset (1)	
7	Unit directory CRC (calculated)	
12 ₁₆	unit_spec_ID (00 609E ₁₆)	
13 ₁₆	unit_sw_version (01 0483 ₁₆)	
38 ₁₆	command_set_spec_ID	
39 ₁₆	command_set	
54 ₁₆	csr_offset (00 4000 ₁₆)	
3A ₁₆	Logical unit characteristics (00 0A08 ₁₆)	
14 ₁₆	Device type and LUN (0)	
		least significant

Figure D-1 – Basic configuration ROM

The ROM CRC in the first quadlet is calculated on the seventeen quadlets of ROM information that follow.

D.1.1 Root directory

The *node_options* field represents a collection of bits and fields specified in 7.3.2. The value shown, 00FF 2000₁₆, represents basic characteristics of a device that is not isochronous capable. This value is composed of a *cyc_clk_acc* field with a value of FF₁₆ and a *max_rec* value of two. The *max_rec* field encodes a maximum payload of eight bytes in block write requests addressed to the target.

The Node_Capabilities entry in the root directory, with a *key* field of 0C₁₆, has a value where the *spt*, *64*, *fix*, *lst* and *drq* bits are all one. This is a minimum requirement for targets.

The Module_Vendor_ID entry in the root directory, with a *key* field of 03₁₆, is immediately followed by a textual descriptor leaf entry, with a *key* field of 81₁₆, whose *indirect_offset* value points to a leaf that contains an ASCII string that identifies the vendor. Although this textual descriptor leaf is not shown, if it were placed immediately after the eighteen quadlets illustrated the value of *indirect_offset* would be eleven. See C.2.2 for an example of a vendor identification text leaf.

The Unit_Directory entry in the root directory, with a *key* field of D1₁₆, has an *indirect_offset* value of one that points to the unit directory.

D.1.2 Unit directory

The Command_Set_Spec_ID and Command_Set entries, with a *key* field of 38₁₆ and 39₁₆, respectively, are expected to define the command set used by the target.

The Management_Agent entry in the unit directory, with a *key* field of 54₁₆, has a *csr_offset* value of 00 4000₁₆ that indicates that the management agent CSR has a base address of FFFF F001 0000₁₆ within the node's initial memory space.

The Logical_Unit_Characteristics entry in the unit directory, with a *key* field of 3A₁₆, has an immediate value of 00 0A08₁₆. This indicates a target that implements the basic task management model, may reorder tasks without restriction and does not support isochronous operations. In addition, the target is expected to complete task management requests (including login) within five seconds and fetches 32-byte ORB's.

The Logical_Unit_Number entry in the unit directory, with a *key* field of 14₁₆, has an immediate value of zero that indicates a device with a logical unit number of zero.

D.2 SCSI command set target

The sample bus information block and root directory for basic targets are equally applicable to targets that use SCSI command set(s) and are not repeated below. Figure D-2 shows an example of a unit directory and textual descriptor leaves for a SCSI direct-access device.

most significant	
9	Unit directory CRC (calculated)
12 ₁₆	unit_spec_ID (00 609E ₁₆)
13 ₁₆	unit_sw_version (01 0483 ₁₆)
38 ₁₆	command_set_spec_ID (00 609E ₁₆)
39 ₁₆	command_set (01 04D8 ₁₆)
54 ₁₆	csr_offset (00 4000 ₁₆)
3A ₁₆	Logical unit characteristics (00 0A08 ₁₆)
14 ₁₆	Device type and LUN (0)
17 ₁₆	model_ID
81 ₁₆	Text leaf offset (5)
3	Text leaf CRC (calculated)
spec_type (0)	specifier_ID (0)
language_ID (0)	
5431 3000 ₁₆ (ASCII "T10")	
3	Text leaf CRC (calculated)
spec_type (0)	specifier_ID (0)
language_ID (0)	
5151 5151 ₁₆ (ASCII "QQQQ")	
least significant	

Figure D-2 – SCSI configuration ROM

D.2.1 Unit_Directory

The Command_Set_Spec_ID and Command_Set_Version entries, with *key* fields of 38₁₆ and 39₁₆, respectively, specify that the target uses one of the SCSI command sets.

The Management_Agent entry in the unit directory, with a *key* field of 54₁₆, has a *csr_offset* value of 00 4000₁₆ that indicates that the MANAGEMENT_AGENT register has a base address of FFFF F001 0000₁₆ within the node's initial memory space.

The Logical_Unit_Characteristics entry in the unit directory, with a *key* field of $3A_{16}$, has an immediate value of $00\ 0A08_{16}$. This indicates a target that implements the basic task management model, may reorder tasks without restriction and does not support isochronous operations. In addition, the target is expected to complete task management requests (including login) within five seconds and fetches 32-byte ORB's.

The Logical_Unit_Number entry in the unit directory, with a *key* field of 14_{16} , has an immediate value of zero; this indicates a direct-access device that uses the SCSI block command (SBC) set and whose logical unit number is zero.

The Model_ID entry in the unit directory, with a *key* field of 17_{16} , has an immediate value whose meaning is specified by the module vendor. Immediately following the Model_ID entry is a textual descriptor leaf entry, with a *key* field of 81_{16} , whose *indirect_offset* value of 5 points to a leaf that contains the ASCII string "QQQQ".

D.2.2 Textual descriptor leaves

Textual descriptor leaf entries, specified by ISO/IEC 13213:1994, permit text strings to be associated with the immediately preceding configuration ROM entry. In this example, two textual descriptor leaves are shown to illustrate a product made by the T10 company with a model identification of QQQQ.

The first textual descriptor leaf is associated with the Module_Vendor_ID entry in the root directory (not shown). The second leaf is associated with the Model_ID entry in the unit directory. The text strings are analogous to the vendor and product identification fields reported in INQUIRY data.

Because textual descriptor leaves are useful to device discovery and management software (in order to display meaningful messages for a user), implementers are encouraged to include textual descriptor leaves for at least the two instances shown.

Annex E (informative)

Serial Bus transaction error recovery

Inherent in the nature of Serial Bus as a split-transaction bus are transaction errors that can leave the requester and responder with different or ambiguous information. One instance occurs when an acknowledge packet transmitted after receipt of a request or response packet is corrupted and not observed by the sender of the primary packet; the same ambiguity exists after a split-transaction time-out.

When an acknowledge packet is missed by the sender of the primary packet, the sender does not know which of the following applies:

- The primary packet was correctly received by the destination node (and resultant side-effects in that node may have occurred); or
- The primary packet had a CRC or other error, has not been correctly received by the destination node and no state changes have occurred.

When a split time-out occurs after a request subaction, the sender knows that the packet was correctly received by the destination node but does not know whether or not resultant side-effects have taken place.

NOTE – Draft standard IEEE P1394a contains important information about split time-out errors, the SPLIT_TIMEOUT register and the different procedures the requester and responder shall follow.

If an acknowledge is missing after transmission of a response packet, IEEE Std 1394-1995 prohibits retransmission of the response packet. Even in the case of a missing acknowledgement following a request packet, it may not be advisable to retry (because of side effects associated with certain SBP-2 transactions).

A few of the more common error scenarios and the recommended error recover for each are described below.

E.1 MANAGEMENT_AGENT write request

When a management ORB is signaled to a target by means of an eight-byte block write to the target's MANAGEMENT_AGENT register and no acknowledgement is received, the initiator does not know whether or not the ORB will be fetched by the target.

Error recovery is straightforward if the initiator waits a minimum of *mgt_ORB_timeout* for the return of a status block before any attempt is made to retry the management ORB. By waiting the specified time the initiator avoids the possibility of multiple status blocks for the same ORB address.

E.2 ORB_POINTER write request

A write to the ORB_POINTER register is valid only when the target fetch agent is into be in the RESET or SUSPENDED state. A consequence of the write is that, if successful, the target fetch agent transitions to the ACTIVE state. If no acknowledgement is received by the initiator after a write to the ORB_POINTER register, the initiator should not retry the write. The recommended method for error recovery is a write to the AGENT_RESET register.

E.3 Data buffer, ORB or page table read request

If the target transmits a block read request and receives no acknowledgement, the read request may be retried immediately but care should be taken to not reuse the same transaction label as the failed request. The target shall wait a minimum of a SPLIT_TIMEOUT period before the transaction label may be reused in any subsequent request subaction.

E.4 Status FIFO write request

When the target detects a missing acknowledgement after a write to an initiator's status FIFO, it shall take no error recovery actions. Any target resources allocated to the ORB should be released by the target. The initiator is expected to discover the error by means of a higher-level mechanism, such as a command time-out and to initiate appropriate error recovery. The nature of the error recovery undertaken by the initiator likely depends whether or not the target processes ORB's and returns their status in order, but in any event is beyond the scope of this description.

Annex F (informative)

SCSI Architecture Model

This annex provides information useful to systems implementers: it relates the facilities provided by SBP-2 to the terminology used by the SCSI Architecture Model.

F.1 Object definitions

The SCSI Architecture Model defines objects within the SCSI domain. The equivalency of SBP-2 objects is enumerated below in those cases where the correlation may not be clear and in those cases where SBP-2 restricts the scope of an object

Initiator identifier: The *login_ID* returned by a target in response to a successful login is the initiator identifier for normal command block requests.

Logical unit number: SBP-2 restricts the scope of the logical unit number to 2^{16} . The *lun* field in management ORB's is one doublet.

NOTE – Neither the *login_ID* nor *lun* fields are present in SCSI command block ORB's, since the value of both is implicit in the CSR addresses of the target fetch agent to which the ORB's are signaled.

Tag: The Serial Bus address of an ORB is the tag by which the task is identified. This requires that initiator memory allocated to a request not be released or reused while the task is active within a task set. The scope of an SBP-2 tag is that of a Serial Bus address, 2^{64} , and equal to the scope defined by SAM-2.

Target identifier: Targets are identified by means of a unit unique ID, or EU-64, found in the target's configuration ROM. When a unit unique ID leaf is not present in configuration ROM, the value of the unit unique ID is construed to be equal to the node unique ID. The scope of a target identifier, 2^{64} , is identical to the scope of a target identifier specified by SAM-2.

Task set: An SBP-2 task set consists of the linked list of normal command block ORB's that are managed by a single target fetch agent. There is no provision in SBP-2 for untagged tasks; an SBP-2 task set always consists of zero or more tagged tasks.

SBP-2 delimits the extent of a task set in a way that is compatible with SAM-2 but that differs from the definition given in SAM-2 of "...a group of tasks *within* a target..." (emphasis added). By way of contrast with SAM-2, an SBP-2 task enters the task set when it is linked into an active request list. The extent of an SBP-2 task set includes all the uncompleted ORB's linked into a request list in initiator memory, not solely the requests already fetched by the target.

Untagged task: SBP-2 does not define untagged tasks.

F.2 Status

SCSI status is reported in the *status* field, which is part of the status block defined in B.2.

F.3 Command delivery services

SAM-2 requires that four protocol services be defined to support the Execute Command remote procedure call. The SBP-2 facilities used to provide these services are described below.

F.3.1 Send SCSI Command

The formal arguments of the Send SCSI Command service are:

Task address,
CDB,
[Task Attribute],
[Data-out buffer],
[Command byte count],
[Autosense request]

The task address argument is composed of the address of the ORB and the logical unit for which it is intended. The logical unit is implicit in the target fetch agent to which the ORB is signaled. The request is signaled to a target fetch agent by the methods specified by section 9.

The CDB argument is encapsulated within an ORB and fetched by the target from initiator memory.

The task attribute argument, either SIMPLE or ORDERED, is implicit in the target implementation. The Logical_Unit_Characteristics entry in the unit directory indicates which task attribute is implemented. When the *ordered* bit in this entry is zero, the all tasks have an implicit attribute of SIMPLE. Otherwise, if *ordered* is set to one, the task attribute is ORDERED for all tasks.

The data-out buffer argument, if present, is specified by the *data_descriptor* and *data_size* fields in the ORB.

SBP-2 provides no means by which the command byte count argument may be communicated to the target. The target may determine the length of the command by an examination of the operation code in the CDB.

The SBP-2 status block provides for the return of autosense data; the initiator may be expected to reformat the information as SCSI sense data before it is presented to the application client.

F.3.2 SCSI Command Received

The formal arguments of the SCSI Command Received service are:

Task address,
[Task Attribute],
CDB,
[Autosense request]

The task address argument is composed of the address of the ORB and the logical unit for which it is intended. The logical unit is implicit in the target fetch agent to which the ORB is signaled. A Serial Bus write indication for the AGENT_RESET register signals the target fetch agent that there may be a new SCSI command. The details of this indication are specified in section 9.

The task attribute argument is implicit in the target implementation and may be determined by an examination of the *ordered* bit in the Logical_Unit_Characteristics entry in configuration ROM.

The CDB argument is encapsulated within the ORB and fetched by the target from initiator memory.

The *status_FIFO* address, provided by the initiator as part of the login procedure, is the address for the return of autosense data.

F.3.3 Send Command Complete

The formal arguments of the Send Command Complete service are:

Task identifier,
[Sense data],
Status,
Service response

The task identifier argument is derived from the address at which the status information is stored. The ORB specified the address for the status block, either implicitly by means of a fixed offset from the address of the ORB or explicitly by means of the *status_FIFO* field. In either case, the initiator ensures that the address at which status is stored is sufficient to uniquely correlate the status with the task identifier.

SCSI sense data may be returned in the status block upon completion of a SCSI command.

The service response argument is encoded within the status block by *resp* and *sbp_status* as summarized below.

Table F-1 – SAM-2 Service responses

Service response	<i>resp</i>	<i>sbp_status</i>	Description
Task complete	0	0	The task has ended with a completion status indicated by <i>status</i>
Linked command complete	—	—	Not supported by SBP-2
Linked command complete (with flag)	—	—	Not supported by SBP-2
Function complete	0	0	Used by task management functions
Service delivery or target failure	1	Various (as defined by SBP-2)	The command has completed because of a Serial Bus service failure or a target malfunction
Function rejected	0	9	Used by task management functions

F.3.4 Command Complete Received

The formal arguments of the Command Complete Received service are:

Task address,
[Data-in buffer],
[Sense data],
Status,
Service response

The task address argument is derived from the address to which the status information was stored. The ORB specified the address for the status block, either implicitly by means of a fixed offset from the address of the ORB or explicitly by means of the *status_FIFO* field. In either case, the initiator ensures that the address at which status is stored is sufficient to uniquely correlate the status with the task identifier.

SCSI sense data may be returned in the status block upon completion of a SCSI command.

The service response argument is encoded within the status block by *resp*, as described by Table F-1.

F.4 Data transfer services

SAM-2 requires that four protocol services be defined to support data transfer necessary for the Execute Command remote procedure call. The SBP-2 facilities used to provide these services are described below.

F.4.1 Send Data-in

The formal arguments of the Send Data-in service are:

- Task identifier,
- Device server buffer,
- Application client buffer offset,
- Request byte count

The task identifier argument is the Serial Bus address of the ORB for the active task. It is expected that target implementations reference a copy of the ORB maintained in the device's local memory, although nothing precludes a fetch of the information from the initiator memory occupied by the ORB.

The device service buffer argument is vendor-dependent. The data available in the device server buffer is formed into Serial Bus write transactions as described below.

The application client buffer offset is a value maintained by the device server to correlate medium locations with locations in the application client buffer. The base of the application client buffer is specified by the *data_descriptor* field supplied by the initiator.

The request byte count is determined by the device server.

The target uses one or more Serial Bus quadlet or block write requests to store the requested data into the application client buffer. For the sake of efficiency, it is expected that the target uses the largest block write requests permitted by the *max_payload* field in the ORB and transmit these requests at the speed mandated by the *spd* field in the ORB.

F.4.2 Data-in Delivered

The formal arguments of the Data-in Delivered service are:

- Task identifier

Upon completion of each of the quadlet or block write requests initiated as a result of Send Data-in, the target receives Serial Bus write response confirmations. It is the target's responsibility to correlate the Serial Bus addresses (for which write responses are received) with the task identifier in order to provide the Data-in Delivered confirmation.

F.4.3 Receive Data-out

The formal arguments of the Receive Data-out service are:

- Task identifier,
- Application client buffer offset,

Request byte count,
Device server buffer

The task identifier argument is the Serial Bus address of the ORB for the active task. It is expected that target implementations reference a copy of the ORB maintained in the device's local memory, although nothing precludes a fetch of the information from the initiator memory occupied by the ORB.

The application client buffer offset is a value maintained by the device server to correlate medium locations with locations in the application client buffer. The base of the application client buffer is specified by the *data_descriptor* field supplied by the initiator.

The request byte count is determined by the device server.

The device service buffer argument is vendor-dependent. The data obtained from Serial Bus read responses is moved to the device server buffer as described below.

The target uses one or more Serial Bus quadlet or block read requests to fetch the requested data from the application client buffer. For the sake of efficiency, it is expected that the target use the largest block read requests permitted by the *max_payload* field in the ORB and transmit these requests at the speed mandated by the *spd* field in the ORB.

F.4.4 Data-out received

The formal arguments of the Data-out received service are:

Task identifier

Upon completion of each of the quadlet or block read requests initiated as a result of Receive Data-out, the target receives Serial Bus read response confirmations and their accompanying data. It is the target's responsibility transfer the data to the device server buffer and to correlate the Serial Bus addresses (for which read responses are received) with the task identifier in order to provide the Data-out Received confirmation.

F.5 Contingent allegiance

SBP-2 targets implement SCSI-2 contingent allegiance and do not support CDB's whose *naca* bit in the control byte is one. The contingent allegiance condition exists within a task set when a logical unit stores a status block for a command where *status* is set to CHECK CONDITION or COMMAND TERMINATED. Since SBP-2 targets implement autosense *via* the return of the status block, the contingent allegiance condition is automatically cleared.

At the time a contingent allegiance condition is created, the logical unit:

- a) immediately halts the operations of the fetch agent for the faulted initiator;
- b) aborts the task set in the same fashion as if an ABORT TASK SET task management function had been signaled to the target;
- c) clears the contingent allegiance condition.

Because the faulted initiator's fetch agent has been halted, it is necessary for the initiator to reset and reinitialize the fetch agent before any commands may be signaled to the target.

F.6 Asynchronous event reporting

SBP-2 does not support asynchronous event reporting as defined by SAM-2.

F.7 Autosense

SBP-2 supports autosense through the return of a status block to the address specified by the login parameter *status_FIFO*. The status block is always stored in the event of an exception condition, *e.g.*, CHECK CONDITION or TERMINATE TASK.

F.8 Hard reset

A Serial Bus reset causes the target to execute a hard reset, as defined by SAM-2.

F.9 Task management functions

SBP-2 targets implement different levels of task management functions, according to the queuing model supported. A target may support a basic or full task set model. The basic and full task management models are as specified by SAM-2.

The relationship between the task management model and the task management functions supported is given by the table below.

Function	Basic	Full	Comments
ABORT TASK	Required	Required	
ABORT TASK SET	Required	Required	May also be performed directly through the AGENT_RESET register
CLEAR ACA	—	—	SBP-2 devices implement SCSI-2 contingent allegiance
CLEAR TASK SETS	Optional	Required	
TARGET RESET	Required	Required	May also be performed directly through the RESET_START register
LOGICAL UNIT RESET	Optional	Optional	Functions as TARGET RESET but scope is limited to a single logical unit
TERMINATE TASK	Not supported	Optional	The basic model provides no control over individual tasks

An initiator may consult the Logical_Unit_Characteristics entry in configuration ROM to determine which task management model is implemented.

SAM-2 additionally requires protocol services to support the task management functions, enumerated below. In all of the definitions for the task management function protocol services, the following apply:

- The object address is as specified by SAM-2 and modified by the SBP-2 object definitions;
- The function identifier is one of the six task management functions defined by SAM-2; and
- The service response is one of Function complete, Function rejected or Service delivery or target failure. These service responses are encoded by the *resp* and *sbp_status* fields in the status block stored by the target upon completion of a request. See Table F-1 for the numeric values that encode the service responses.

F.9.1 Send task management request

The formal arguments of the Send task management request service are:

Object address,
Function identifier

Subsequent to the creation of a task management ORB in initiator memory, the initiator signals the request to the target management agent by the methods described in SBP-2.

F.9.2 Task management request received

The formal arguments of the Task management request received service are:

Object identifier,
Function identifier

When a Serial Bus write indication is received for the target's MANAGEMENT_AGENT register, the target may fetch the request from initiator memory.

F.9.3 Task management function executed

The formal arguments of the Task management function executed service are:

Object identifier,
Service response

The target signals the completion of the task management function by storing an 8-byte status block at the address specified by *status_FIFO* in the ORB.

F.9.4 Received Function-executed

The formal arguments of the Received Function-executed service are:

Object address,
Service response

When the initiator receives a Serial Bus write indication for data addressed to the *status_FIFO* address, it may examine the status block to determine the service response from *resp*.

Annex G (informative)

AV/C Encapsulation

Devices that use the AV/C Digital Interface Command Set use IEC 61883/FDIS Function Control Protocol (FCP) for the encapsulation of command and status. This annex illustrates how SBP-2 could be utilized as the transport layer in place of FCP.

AV/C devices are consumer electronic devices such as camcorders, VCR's, stereo tuners and televisions—this is not an exhaustive list. Their commands, status and operations are standardized by the 1394 Trade Association Specification for AV/C Digital Interface Command Set, Version 2.0, April 22, 1977. These devices present or accept most of their data over Serial Bus isochronous channels; asynchronous requests are used for control information.

AV/C devices differ from the isochronous model presented in 4.7 in two important ways: *a)* there is a single command queue and *b)* isochronous flow control is accomplished by plug control registers (PCR's) and commands contained within the single queue. These differences are illustrative, not only for AV/C devices but also for other, future devices which may use a similar model.

This annex is particular but the intent is general: by way of an AV/C example, it provides guidance for the use of SBP-2 for any device that follows an equivalent, single-queue model.

G.1 Logical unit, unit and subunit models

SBP-2 describes a hierarchical device structure composed of units with subordinate logical units. AV/C has a hierarchical structure of units with subordinate subunits.

An AV/C unit is described by configuration ROM entries in a unit directory and is implemented as a single logical unit. AV/C subunits are not visible as SBP-2 objects; a single queue is maintained for commands for all the subunits of an AV/C unit.

G.2 AV/C command frame

An AV/C command frame is a variable-length data structure between four and 512 bytes in length. A normal command block ORB can accommodate 12 bytes of command frame, as illustrated below.

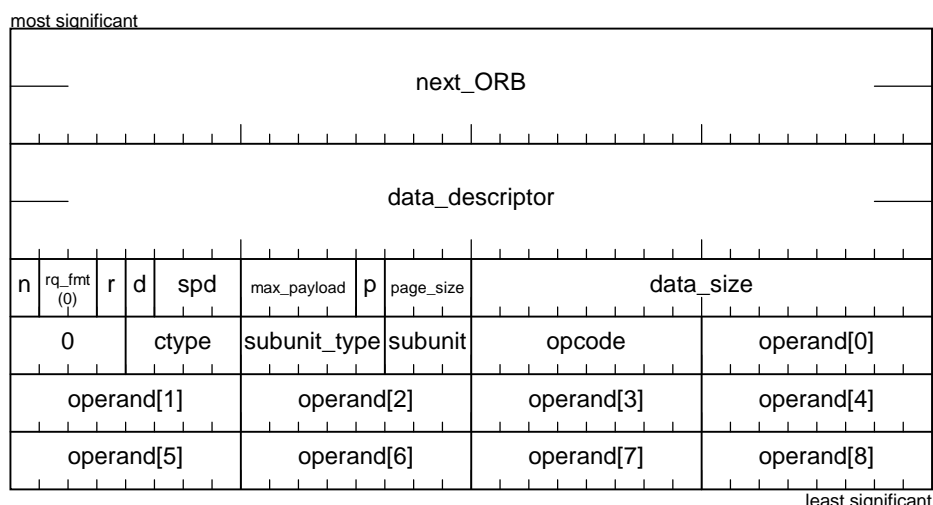


Figure G-1 – AV/C command frame ORB

The *next_ORB*, *data_descriptor*, *rq_fmt*, *spd*, *max_payload*, *page_size* and *data_size* fields and the *notify*, *direction* and *page_table_present* bits (abbreviated as *n*, *d*, and *p*, respectively, in the figure above) are as specified in 5.1.2.1. The AV/C commands described in the Version 2.0 1394 Trade Association specification do not reference a data buffer; consequently the *data_size* field is always zero in an AV/C command frame ORB.

The *ctype*, *subunit_type*, *subunit*, *opcode* and *operand* fields are specified by AV/C.

G.3 AV/C response frame

An AV/C response frame is a variable-length data structure between four and 512 bytes in length. The status block could hold up to 24 bytes of response frame, but since this exceeds the capacity of the AV/C command ORB the size of the status block is restricted to 20 bytes, as illustrated below.

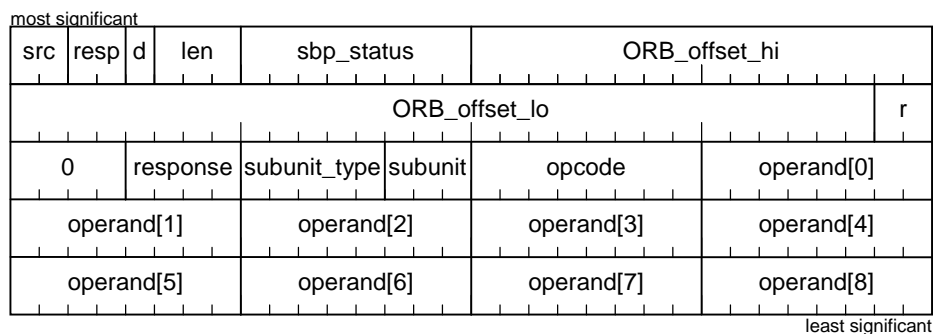


Figure G-2 – Status block format for AV/C response frame

The *src*, *resp*, *sbp_status* fields and the *dead* bit (abbreviated as *d* in the figure above) are as specified in 5.3.

The *len* field indicates the size of the status block and shall have a value of four.

The *ORB_offset_hi* and *ORB_offset_lo* fields together identify the AV/C command frame to which the AV/C response frame pertains. This provides positive command and response matching—an enhanced capability in comparison to FCP.

The *response*, *subunit_type*, *subunit*, *opcode* and *operand* fields are specified by AV/C.

G.4 Configuration ROM

The sample bus information block and root directory for basic targets illustrated in Annex D are equally applicable to targets that use the AV/C command set and are not repeated below. Figure G-3 shows an example of a unit directory and textual descriptor leaves for an AV/C device.

most significant															
7								Unit directory CRC (calculated)							
12 ₁₆								unit_spec_ID (00 609E ₁₆)							
13 ₁₆								unit_sw_version (01 0483 ₁₆)							
38 ₁₆								command_set_spec_ID (00 A02D ₁₆)							
39 ₁₆								command_set (01 0001 ₁₆)							
54 ₁₆								csr_offset (00 4000 ₁₆)							
3A ₁₆								Logical unit characteristics (40 0A08 ₁₆)							
14 ₁₆								Unit type and number (1F 0000 ₁₆)							
least significant															

Figure G-3 – AV/C unit directory

The Command_Set_Spec_ID and Command_Set_Version entries, with a *key* field of 38₁₆ and 39₁₆, respectively, specify that the target uses the AV/C command set defined by the Version 2.0 1394 Trade Association Specification.

The Management_Agent entry in the unit directory, with a *key* field of 54₁₆, has a *csr_offset* value of 00 4000₁₆ that indicates that the MANAGEMENT_AGENT register has a base address of FFFF F001 0000₁₆ within the node's initial memory space.

The Logical_Unit_Characteristics entry in the unit directory, with a *key* field of 3A₁₆, has an immediate value of 40 0A08₁₆. This indicates a target that implements the basic task management model, executes commands in the order queued and does not support (the dual-queue model of) isochronous operations. In addition, the target is expected to complete a login within five seconds and fetches 32-byte ORB's.

The Logical_Unit_Number entry in the unit directory, with a *key* field of 14₁₆, has an immediate value of 1F 0000₁₆; this identifies logical unit zero and indicates that AV/C commands are used to query the number and type of subunits associated with the unit.

G.5 Operations

Control of an AV/C device implemented with SBP-2 follows the steps described in sections 8 and 9. Subsequent to a successful login, the controller may create a queue of one or more ORB's that hold AV/C

command frames and signal these to the target. As the commands complete, the response frames are returned in status blocks addressed to the *status_FIFO* specified by the controller at login.

SBP-2 offers a number of advantages over FCP for the delivery of AV/C command and response frames:

- The target may pace command processing to suit its own requirements and resource availability. Consequently the controller should not expect busy responses and does not have to provide error recovery for busy conditions;
- The correlation between a particular AV/C command frame and its response frame is explicit and unambiguous—even in the case of complex AV/C devices that may be able to queue or process more than one command at a time;
- SBP-2 permits the controller to manage the commands in progress at the target and be certain of the completion status of pending commands if one or more commands are terminated prematurely; and
- The AV/C command frame ORB has descriptor fields for a data buffer that permit the target to use a "pull" model for asynchronous data transfer to or from the controller. The target paces data delivery between itself and the controller; FCP has no such model.

As AV/C command sets are developed for newer and more sophisticated devices, these SBP-2 features may become increasingly important to systems implementers.