

Information technology - SCSI Architecture Model - 3 (SAM-3)

This is an internal working document of T10, a Technical Committee of Accredited Standards Committee INCITS (InterNational Committee for Information Technology Standards). As such this is not a completed standard and has not been approved. The contents may be modified by the T10 Technical Committee. The contents are actively being modified by T10. This document is made available for review and comment only.

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

T10 Technical Editor:

Ralph O. Weber
ENDL Texas
18484 Preston Road
Suite 102 PMB 178
Dallas, TX 75252
USA

Telephone: 214-912-1373
Facsimile: 972-596-2775
Email: ROWeber@IEEE.org

Reference number
ISO/IEC 14776-413-200x
ANSI INCITS.***:200x

Points of Contact:

T10 Chair

John B. Lohmeyer
LSI Logic
4420 Arrows West Drive
Colorado Springs, CO 80907-3444
Tel: (719) 533-7560
Fax: (719) 533-7183
Email: lohmeier@t10.org

T10 Vice-Chair

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

INCITS Secretariat

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

Telephone: 202-737-8888
Facsimile: 202-638-4922
Email: INCITS@itic.org

T10 Web Site www.t10.org

T10 Reflector To subscribe send e-mail to majordomo@T10.org with 'subscribe' in message body
To unsubscribe send e-mail to majordomo@T10.org with 'unsubscribe' in message body
Internet address for distribution via T10 reflector: T10@T10.org

Document Distribution

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

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

or

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

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

Draft

**American National Standards
for Information Systems -**

SCSI Architecture Model - 3 (SAM-3)

Secretariat
InterNational Committee for Information Technology Standards

Approved mm dd yy

American National Standards Institute, Inc.

Abstract

This standard specifies the SCSI Architecture Model. The purpose of the architecture is to provide a common basis for the coordination of SCSI standards and to specify those aspects of SCSI I/O system behavior that are independent of a particular technology and common to all implementations.

Draft

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 toward their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he or she 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 in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

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

As of the date of publication of this standard and following calls for the identification of patents that may be required for the implementation of the standard, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Published by
American National Standards Institute
11 West 42nd Street, New York, NY 10036

Copyright 3/17/04 by American National Standards Institute
All rights reserved.

Printed in the United States of America

Draft

Contents

	Page
1 Scope	1
1.1 Introduction	1
1.2 Requirements precedence	2
1.3 SCSI standards family	3
2 Normative references	6
2.1 Normative references	6
2.2 Approved references	6
2.3 References under development	6
3 Definitions, symbols, abbreviations, and conventions	7
3.1 Definitions	7
3.2 Acronyms	14
3.3 Keywords	14
3.4 Editorial conventions	15
3.5 Numeric conventions	16
3.6 Notation conventions	16
3.6.1 Hierarchy diagram conventions	16
3.6.2 Notation for procedure calls	17
3.6.3 Notation for state diagrams	18
4 SCSI architecture model	19
4.1 Introduction	19
4.2 The SCSI distributed service model	20
4.3 The SCSI client-server model	21
4.4 The SCSI structural model	22
4.5 SCSI domain	24
4.6 The service delivery subsystem	24
4.6.1 The service delivery subsystem object	24
4.6.2 Synchronizing client and server states	25
4.6.3 Request/Response ordering	25
4.7 SCSI devices	26
4.7.1 SCSI initiator device	27
4.7.2 SCSI target device	28
4.7.3 SCSI target/initiator device	29
4.7.4 SCSI port identifier	30
4.7.5 Relative port identifier	30
4.7.6 SCSI task router	30
4.7.7 SCSI device name	31
4.7.8 SCSI port name	31
4.8 Logical units	32
4.9 Logical unit numbers	33
4.9.1 Logical unit numbers overview	33
4.9.2 LUN 0 address	33
4.9.3 Single level logical unit number structure	33
4.9.4 Eight byte logical unit number structure	35
4.9.5 Logical unit addressing method	37
4.9.6 Peripheral device addressing method	37
4.9.7 Flat space addressing method	38
4.9.8 Extended logical unit addressing	39
4.9.9 Well known logical unit addressing	41

4.9.10 Logical unit not specified addressing.....	41
4.10 Well known logical units	42
4.11 Tasks and task tags.....	42
4.12 The nexus object	43
4.13 SCSI ports	44
4.13.1 SCSI port configurations.....	44
4.13.2 SCSI devices with multiple ports	44
4.13.3 Multiple port target SCSI device structure	45
4.13.4 Multiple port initiator SCSI device structure.....	46
4.13.5 Multiple port target/initiator SCSI device structure	47
4.13.6 SCSI initiator device view of a multiple port SCSI target device	48
4.13.7 SCSI target device view of a multiple port SCSI initiator device	50
4.14 Model for dependent logical units.....	51
4.15 The SCSI model for distributed communications	53
5 SCSI command model	59
5.1 The Execute Command procedure call	59
5.2 Command descriptor block (CDB).....	61
5.3 Status	62
5.3.1 Status codes.....	62
5.3.2 Status precedence.....	64
5.4 SCSI transport protocol services in support of Execute Command.....	65
5.4.1 Overview.....	65
5.4.2 Execute Command request/confirmation SCSI transport protocol services.....	65
5.4.3 Data transfer SCSI transport protocol services	67
5.4.3.1 Introduction.....	67
5.4.3.2 Data-In delivery service	69
5.4.3.3 Data-Out delivery service	69
5.4.3.4 Terminate Data Transfer service	70
5.5 Task and command lifetimes.....	71
5.6 Task management function lifetime	72
5.7 Aborting tasks.....	72
5.7.1 Mechanisms that cause tasks to be aborted	72
5.7.2 When a SCSI initiator port aborts its own tasks	73
5.7.3 When a SCSI initiator port aborts tasks from other SCSI initiator ports	73
5.8 Command processing examples	74
5.8.1 Unlinked command example	74
5.8.2 Linked command example.....	75
5.9 Command processing considerations and exception conditions.....	76
5.9.1 Commands that complete with CHECK CONDITION status.....	76
5.9.1.1 Overview.....	76
5.9.1.2 Handling tasks when ACA is not in effect.....	76
5.9.1.3 Aborting other tasks when CHECK CONDITION status is returned without establishing an ACA.....	77
5.9.2 Auto contingent allegiance (ACA).....	77
5.9.2.1 Establishing an ACA.....	77
5.9.2.2 Handling new tasks from the faulted initiator port when ACA is in effect	79
5.9.2.3 Handling new tasks from non-faulted initiator ports when ACA is in effect	79
5.9.2.3.1 Commands permitted from non-faulted initiator ports during ACA	79
5.9.2.3.2 Handling new tasks from non-faulted initiator ports when ACA is in effect	80
5.9.2.4 Clearing an ACA condition	80
5.9.3 Overlapped commands	81
5.9.4 Incorrect logical unit selection	81
5.9.5 Task attribute exception conditions	82
5.9.6 Sense data	82

5.9.7 Unit Attention condition.....	82
6 SCSI events and event notification model	84
6.1 SCSI events overview	84
6.2 Establishing a unit attention condition subsequent to detection of an event	85
6.3 Conditions resulting from SCSI events.....	86
6.3.1 Power on	86
6.3.2 Hard reset.....	86
6.3.3 Logical unit reset	87
6.3.4 I_T nexus loss	87
6.4 Event notification SCSI transport protocol services.....	88
7 Task management functions	89
7.1 Introduction.....	89
7.2 ABORT TASK.....	90
7.3 ABORT TASK SET.....	90
7.4 CLEAR ACA	91
7.5 CLEAR TASK SET	91
7.6 LOGICAL UNIT RESET.....	92
7.7 QUERY TASK	92
7.8 Task management SCSI transport protocol services	92
7.9 Task management function example.....	95
8 Task set management	96
8.1 Introduction to task set management	96
8.2 Implicit head of queue	96
8.3 Task management models	96
8.3.1 Task management model management features	96
8.3.2 Full task management model	96
8.3.3 Basic task management model	97
8.4 Task management events	97
8.5 Task states	98
8.5.1 Overview.....	98
8.5.1.1 Task state nomenclature	98
8.5.1.2 Suspended information.....	98
8.5.2 Enabled task state	98
8.5.3 Blocked task state	98
8.5.4 Dormant task state	98
8.5.5 Ended task state.....	98
8.5.6 Task states and task lifetimes	99
8.6 Task attributes	100
8.6.1 Overview.....	100
8.6.2 Simple task.....	100
8.6.3 Ordered task.....	100
8.6.4 Head of queue task	100
8.6.5 ACA task.....	100
8.7 Task priority	101
8.8 Task state transitions.....	102
8.9 Task set management examples.....	103
8.9.1 Introduction.....	103
8.9.2 Head of queue tasks.....	104
8.9.3 Ordered tasks	106
8.9.4 ACA task.....	107

Annex A

Identifiers and names for objects 108

 A.1 Identifiers and names overview..... 108

 A.2 Identifiers and names..... 108

 A.3 SCSI transport protocol acronyms and bibliography 112

Annex B

Terminology mapping 114

Tables

	Page
1 Single level logical unit number structure for logical unit numbers 255 and below	33
2 Single level logical unit number structure for logical unit numbers 16 383 and below	34
3 Eight byte logical unit number structure adjustments	35
4 Eight Byte logical unit number structure	36
5 Format of addressing fields.....	36
6 ADDRESS METHOD field values.....	37
7 Logical unit addressing	37
8 Peripheral device addressing.....	38
9 Flat space addressing.....	39
10 Extended logical unit addressing	39
11 LENGTH field values	39
12 Two byte extended logical unit addressing format.....	40
13 Four byte extended logical unit addressing format	40
14 Six byte extended logical unit addressing format.....	40
15 Eight byte extended logical unit addressing format	40
16 Logical unit extended address methods	40
17 Well known logical unit extended address format.....	41
18 Logical unit not specified extended address method format.....	41
19 Mapping nexus to SAM-2 identifiers	43
20 CONTROL byte	61
21 Status codes	62
22 Task handling when ACA is not in effect	76
23 Aborting tasks when an ACA is not established	77
24 Blocking and aborting tasks when an ACA is established	78
25 Handling for new tasks from a faulted initiator port during ACA.....	79
26 Handling for new tasks from non-faulted initiator ports during ACA.....	80
27 Unit attention additional sense codes for events detected by SCSI target devices	85
28 Task Management Functions.....	89
29 Task State Nomenclature	98
30 Task attributes	100
31 Task attribute and state indications in examples	104
32 Dormant task blocking boundary requirements	106
A.1 Object size and support requirements.....	108
A.2 Object identifier size for each SCSI transport protocol.....	109
A.3 Object identifier format for each SCSI transport protocol	110
A.4 Object name size for each SCSI transport protocol	111
A.5 Object name format for each SCSI transport protocol.....	112
B.1 SAM-3 to SAM terminology mapping	114

Figures

	Page
1 Requirements precedence	2
2 SCSI document structure	3
3 Example hierarchy diagram	16
4 Example state diagram	18
6 SCSI client-server model	21
7 SCSI I/O system and domain model	22
8 Overall SCSI domain model	23
9 SCSI domain model	24
10 Service delivery subsystem model	24
11 SCSI initiator device model	27
12 SCSI target device model	28
13 SCSI target/initiator device with SCSI target/initiator ports model	29
14 SCSI target/initiator device without SCSI target/initiator ports model	29
15 Logical unit model	32
16 Eight Byte logical unit number structure adjustments	35
17 SCSI device functional models	44
18 Multiple port target SCSI device structure model	45
19 Multiple port SCSI initiator device structure model	46
20 Multiple port target/initiator SCSI device structure model	47
21 SCSI target device configured in a single SCSI domain	48
22 SCSI target device configured in multiple SCSI domains	49
23 SCSI target device and SCSI initiator device configured in a single SCSI domain	49
24 Dependent logical unit model	51
25 Example of hierarchical system diagram	52
26 SCSI transport protocol service model	55
27 Request-Response SAL transaction and related STPL services	55
28 SCSI transport protocol service model for data transfers	56
29 Device server data transfer transaction and related STPL services	57
30 SCSI transport protocol service model for Terminate Data Transfer	57
31 Device server Terminate Data Transfer transaction and related STPL services	58
32 Model for Data-In and Data-Out data transfers	67
33 Command processing events	74
34 Linked command processing events	75
35 Events and event notifications for SCSI target devices	84
36 Events and event notifications for SCSI initiator devices	85
37 Task management processing events	95
38 Example of Dormant state task behavior	99
39 Task states	102
40 Head of queue tasks and blocking boundaries (example 1)	104
41 Head of queue tasks and blocking boundaries (example 2)	105
42 Ordered tasks and blocking boundaries	106
43 ACA task example	107

Revision Information

1 Approved Documents Included

The following T10 approved proposals have been incorporated SAM-3 up to and including this revision:

02-155r6	Response to T10 Letter Ballot comments on SAM-2
02-227r1	Rewrite of SAM-2 example for dependent logical units to remove SCC-isms
02-232r2	Clearing effects of I_T nexus loss
02-241r2	Making Application Clients and Device Servers True Peers in SAM-2
02-244r0	SAM to SAM-2 Changed Terms Annex
02-245r0	SAM-3 Identifiers and names for SAS SSP
02-273	Minutes of T10 Meeting #50 — Change AER protocol requirement from 'shall' to 'may'
02-278r0	SAM-3 Query Task
02-321r0	LUN value meaning unspecified LU
02-348r0	SAM-3 Names and identifiers for iSCSI revision 16
02-366r0	Proposal for adding new arguments for Data-In Delivered and Data-Out Received confirmations into SAM-3
02-403r2	Maximum and first burst size and protocol services
02-404r2	Sense Data Length for Execute Command
02-419r7	Device Identifiers and VPD Data
02-427r1	SAS Target and Initiator Names
02-458	Minutes of T10 Meeting #52 — Remove asynchronous event reporting
02-464r3	SAM-3, SPC-3, and SBC-2 power conditions proposal
02-466r0	SAM-3 definitions for SCSI initiator device and SCSI target device
02-473r2	Hard reset and multiple port devices
02-488r1	SAM-3 Logical Unit Names
03-001r1	Remove TARGET RESET from SAM-3 and SPC-3
03-002r3	Remove SPI from SAM-3 and SPC-3
03-003r2	Remove ULP and LLP from SAM-3
03-004r2	SAM-3 Glossary Cleanup
03-005r1	What is architectural about SAM-3 CDB definitions?
03-032r1	Single level LUN formats
03-108r2	Calls and Arguments in SAM-3
03-173r0	Task Management Incorrect Logical Unit Selection
03-184r1	SAM-3 Hard reset generates I_T nexus loss
03-243r2	Optional INQUIRY command processing before entering task set
03-244r2	New Terminate Data Transfer transport protocol service and related clarification
03-271r2	Obsolete untagged tasks
03-302r4	SAM-3 Per-Command Priority Tagging
03-344r6	Report all initiator and target ports
03-351r2	Reporting task attribute support
03-352r0	Making the DQUE bit obsolete in SAM-3
03-368r2	Status precedence
03-376r4	Invalid LUN on Request Sense Command
04-070r1	Status Precedence Problem
04-087r1	QErr Wording

To the best of the technical editor's knowledge, all proposals for SAM-3 approved by T10 have been included in this revision.

2 Revision History

2.1 Revision 0 (12 July 2002)

SAM-3 revision 0 is exactly SAM-2 revision 23 (i.e., the revision of SAM-2 on which T10 conducted a Letter Ballot in April 2002), except that all change bars have been removed. This prepares SAM-3 to assist T10 in following SAM changes resulting from the Letter Ballot, because SAM-3 revision 1 can include those changes with change bars, an aid that will not be placed in the Public Review revision of SAM-2.

2.2 Revision 1 (22 July 2002)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

02-245r0	SAM-3 Identifiers and names for SAS SSP
02-273	Minutes of T10 Meeting #50 — Change AER protocol requirement from 'shall' to 'may'
02-278r0	SAM-3 Query Task

No other changes were made in revision 1.

2.3 Revision 2 (11 September 2002)

The following proposals were incorporated SAM-3 in this revision because the CAP working group approved them as part of the SAM-2 letter ballot resolution:

02-155r6	Response to T10 Letter Ballot comments on SAM-2
02-227r1	Rewrite of SAM-2 example for dependent logical units to remove SCC-isms
02-241r2	Making Application Clients and Device Servers True Peers in SAM-2
02-244r0	SAM to SAM-2 Changed Terms Annex
02-348r0	SAM-3 Names and identifiers for iSCSI revision 16

Revision 2 of SAM-3 incorporates all the CAP approved changes in response to the SAM-2 letter ballot. Because of coordination issues, these changes have not been approved by T10 as of the time SAM-3 revision 2 is being prepared. If changes are not approved by T10, then they will be removed in revision 3.

2.4 Revision 3 (12 September 2002)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

02-232r2	Clearing effects of I_T nexus loss
----------	------------------------------------

The new clause introduced by 02-232r2 was edited extensively, starting with the clause title and including a reorganization of the information approved for inclusion in the new clause.

Also the capitalization in the clause headings for clause 4 and clause 5 was changed to match that found in the remainder of the working draft.

2.5 Revision 4 (7 November 2002)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 02-321r0 LUN value meaning unspecified LU
- 02-366r0 Proposal for adding new arguments for Data-In Delivered and Data-Out Received confirmations into SAM-3
- 02-427r1 SAS Target and Initiator Names
- 02-458 Minutes of T10 Meeting #52 — Remove asynchronous event reporting
- 02-466r0 SAM-3 definitions for SCSI initiator device and SCSI target device

The November CAP working group reviewed the results of incorporating 02-232r2 in the SCSI events and event notification model clause. Several changes were recommended and those changes have been incorporated in revision 4.

In the glossary, the alphabetization of 'SCSI target...' and 'SCSI transport...' was corrected.

2.6 Revision 5 (12 February 2003)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 02-403r2 Maximum and first burst size and protocol services
- 02-404r2 Sense Data Length for Execute Command
- 02-473r2 Hard reset and multiple port devices
- 02-488r1 SAM-3 Logical Unit Names
- 03-002r3 Remove SPI from SAM-3 and SPC-3
- 03-032r1 Single level LUN formats

Concerns expressed on the T10 reflector regarding change 20 in 03-002r3 prompted the editor to not make the change and to insert an editor's note regarding the text for which changes were proposed. The intent is that a future proposal will resolve the issues raised by 03-002r3 change 20.

In 5.3.2 (Status precedence), one of the two occurrences of RESERVATION CONFLICT was removed.

In the glossary definitions of SCSI initiator port and SCSI target port, terminology usage was changed to match that found in 4.14 (The SCSI model for distributed communications).

The introduction to the task management function prototype was modified to have a structure similar to that approved during the SAM-2 letter ballot for the introduction to the Execute Command function prototype.

In the last paragraph of 7.10 (Task management SCSI transport protocol), QUERY TASK was added to the list of task management functions that include a tag.

In 7.10 (Task management SCSI transport protocol), the function succeeded Service Response was replicated from the task management function prototype definition.

In 5.4.3.2 (Data-In delivery service) and 5.4.3.3 (Data-Out delivery service), the to/from direction identifiers were switched on the descriptions for the Device Server Buffer parameter and text was added to clarify why the transfer direction indicated in these two sentences is backwards from all other SCSI discussions of transfer direction (i.e., in this particular case the description is from the target's point of view, instead of from the initiator's point of view).

2.7 Revision 6 (19 March 2003)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 03-001r1 Remove TARGET RESET from SAM-3 and SPC-3
- 03-004r2 SAM-3 Glossary Cleanup
- 03-005r1 What is architectural about SAM-3 CDB definitions?

Several incorrect uses of the W-LUN acronym were changed to 'well known logical unit'.

One change overlooked by 03-002r3 concerns logical unit numbers. In the parallel SCSI bus, logical unit numbers could contain less than 64 bits. All SCSI transport protocols except SPI provide for 64 bit logical unit numbers. In keeping with the ground work established by 03-002r3, phrasing that allowed logical unit numbers to contain fewer than 64 bits has been removed in this revision.

The front matter describing planned future changes was removed because proposals under consideration or already approved by T10 address all the concerns raised in that material.

The concepts of an 'I_T_L_Q nexus transaction' included in SPC-3 r12 as the result of an editing review meeting was duplicated in SAM-3 for consistency and completeness.

2.8 Revision 7 (9 May 2003)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 02-464r3 SAM-3, SPC-3, and SBC-2 power conditions proposal
- 03-003r2 Remove ULP and LLP from SAM-3
- 03-108r2 Calls and Arguments in SAM-3
- 03-173r0 Task Management Incorrect Logical Unit Selection
- 03-184r1 SAM-3 Hard reset generates I_T nexus loss

The subclause describing hard resets was enhanced to clarify that, although a hard reset on one target port in a given target device shall not affect another target port in the same target device, the effects of the logical unit reset required by the hard reset may cause ancillary effects that are observable on another target port.

The T10 cover page was cleaned up. Usage of the phrase 'SCSI rules' was eliminated. An instance of 'LOGICAL RESET' was corrected to match the task management function name 'LOGICAL UNIT RESET'.

2.9 Revision 8 (11 July 2003)

T10 did not approve any proposals for inclusion in SAM-3 during its July meeting; however, the following non-substantive changes were felt to be important enough to prepare a new revision.

In Aborting other tasks when CHECK CONDITION status is returned without establishing an ACA, the first sentence in the table row for QERR=11b was corrected to apply on to the faulted initiator. In Clearing an ACA condition, an ACA (task attribute) was corrected to ACA.

2.10 Revision 9 (12 September 2003)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 03-243r2 Optional INQUIRY command processing before entering task set
- 03-244r2 New Terminate Data Transfer transport protocol service and related clarification

To support 03-243r2 glossary entries were added for 'command standard' and 'implicit head of queue'.

The function prototype notation for IN, OUT, and end of prototype (double closing parentheses) was made consistent throughout the working draft.

2.11 Revision 10 (26 November 2003)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 02-419r7 Device Identifiers and VPD Data
- 03-271r2 Obsolete untagged tasks
- 03-352r0 Making the DQUE bit obsolete in SAM-3

While incorporating 02-419r7, it was noted that some of the proposed changes to 4.10 (Well known logical units) removed statements allowing multiple target device names. In view of the history of the proposal, the removal of these statements seemed to be contrary to the committee's intent. So the statements were not removed.

A few editorial changes in 03-271r2 that were not germane to the stated purpose of the proposal were ignored.

Support for SPI-5 was explicitly listed as a SAM-2 feature made obsolete in SAM-3. References to SPI-4 were changed to SPI-5 in support of the information provided in Annex A.

In 5.9.2 [Auto contingent allegiance (ACA)], a redundant phrase that probably resulted from an incorrect cut and paste operation was removed from the last sentence in the second paragraph. In 6.2 [Establishing a unit attention condition subsequent to detection of an event], two typographical errors were corrected in the paragraphs following the note.

2.12 Revision 11 (23 January 2004)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

- 03-351r2 Reporting task attribute support
- 03-368r2 Status precedence

Some of the text in 03-351r2 not specifically marked for change differed from the referenced revision of SAM-3. Differences in unmarked text were not always incorporated. The 03-351r2 rewrite of the basic task management model failed to reference the SCSI transport protocol constraints and turned the discussion of the QUEUE ALGORITHM MODIFIER field in the Control mode page into gibberish, so that text was substantially revised. The gap between requiring application clients to specify a task attribute and transport protocols that do not transport the value was closed by stating that such transport protocols shall ignore the value provided by the application client.

Based on the CAP Working Group request to remove discussion of the TRANSCEIVER MODE CHANGED TO LVD and TRANSCEIVER MODE CHANGED TO SINGLE-ENDED additional sense codes from 03-368r2, discussion of the two additional sense codes was removed throughout SAM-3.

2.13 Revision 12 (16 March 2004)

The following T10 approved proposals were incorporated in SAM-3 in this revision:

03-302r4	SAM-3 Per-Command Priority Tagging
03-344r6	Report all initiator and target ports
03-376r4	Invalid LUN on Request Sense Command
04-070r1	Status Precedence Problem
04-087r1	QErr Wording

Instances of 'that specified' were clarified to 'the value specified'.

Editing errors made during the incorporation of the 03-351r2 [Reporting task attribute support] definition of the basic task management model were corrected. Specifically, the definition of basic task management model support for the SIMPLE task attribute was corrected to reflect reality.

Several instances of 'initiator port' and 'target port' (both terms that do not appear in the glossary) were changed to the glossary defined terms 'SCSI initiator port' and 'SCSI target port', respectively. In 6.3.2 [Hard reset], an instance of 'target' was changed to 'SCSI target device'. In 6.3.4 [L_T nexus loss], an ambiguous 'that' was clarified.

2.14 Revision 13 (17 March 2004)

The editorial changes agreed by the 11 and 12 March 2004 editing meeting were incorporated. This was done in a separate revision so that the change bars resulting from the editing meeting could be easily identified.

An error made during the incorporation of 04-087r1 [QErr Wording] regarding tasks aborted by the actions of other initiators was corrected.

Foreword

This foreword is not part of American National Standard INCITS.***:200x.

The purpose of this standard is to provide a basis for the coordination of SCSI standards development and to define requirements, common to all SCSI technologies and implementations that are essential for compatibility with host SCSI application software and device-resident firmware across all SCSI transport protocols. These requirements are defined through a reference model that specifies the behavior and abstract structure that is generic to all SCSI I/O system implementations.

With any technical document there may arise questions of interpretation as new products are implemented. INCITS has established procedures to issue technical opinions concerning the standards developed by INCITS. These procedures may result in SCSI Technical Information Bulletins being published by INCITS.

These Bulletins, while reflecting the opinion of the Technical Committee that developed the standard, are intended solely as supplementary information to other users of the standard. This standard, ANSI INCITS.***:200x, as approved through the publication and voting procedures of the American National Standards Institute, is not altered by these bulletins. Any subsequent revision to this standard may or may not reflect the contents of these Technical Information Bulletins.

Current INCITS practice is to make Technical Information Bulletins available through:

INCITS Online Store	http://www.techstreet.com/INCITS.html
managed by Techstreet	Telephone: 1-734-302-7801 or
1327 Jones Drive	1-800-699-9277
Ann Arbor, MI 48105	Facsimile: 1-734-302-7811

or

Global Engineering	http://global.ihs.com/
15 Inverness Way East	Telephone: 1-303-792-2181 or
Englewood, CO 80112-5704	1-800-854-7179
	Facsimile: 1-303-792-2192

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

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

<<Insert INCITS member list>>

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

<<Insert T10 member list>>

Introduction

The SCSI Architecture Model - 3 (SAM-3) standard is divided into the following clauses and annexes:

Clause 1 is the scope.

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

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

Clause 4 describes the overall SCSI architectural model.

Clause 5 describes the SCSI command model element of the SCSI architecture.

Clause 6 describes the events that may be detected by a SCSI device.

Clause 7 describes the task management functions common to SCSI devices.

Clause 8 describes the task set management capabilities common to SCSI devices.

Annex A summarizes the identifier and name definitions of the SCSI transport protocols.

Annex B identifies differences between the terminology used in this standard and previous versions of this standard.

American National Standard

INCITS.*:200x****American National Standard
for Information Technology -****SCSI Architecture Model - 3 (SAM-3)****1 Scope****1.1 Introduction**

The set of SCSI (Small Computer System Interface) standards consists of this standard and the SCSI implementation standards described in 1.3. This standard defines a reference model that specifies common behaviors for SCSI devices, and an abstract structure that is generic to all SCSI I/O system implementations.

The set of SCSI standards specifies the interfaces, functions, and operations necessary to ensure interoperability between conforming SCSI implementations. This standard is a functional description. Conforming implementations may employ any design technique that does not violate interoperability.

The following architecture model concepts from previous versions of this standard are made obsolete by this standard:

- a) Support for the SPI-5 SCSI transport protocol (except for informational listings in Annex A);
- b) Contingent Allegiance;
- c) The TARGET RESET task management function; and
- d) Untagged tasks.

1.2 Requirements precedence

This standard defines generic requirements that pertain to SCSI implementation standards and implementation requirements. An implementation requirement specifies behavior in terms of measurable or observable parameters that apply to an implementation. Examples of implementation requirements defined in this document are the command descriptor block format and the status values to be returned upon command completion.

Generic requirements are transformed to implementation requirements by an implementation standard. An example of a generic requirement is the hard reset behavior specified in 6.3.2.

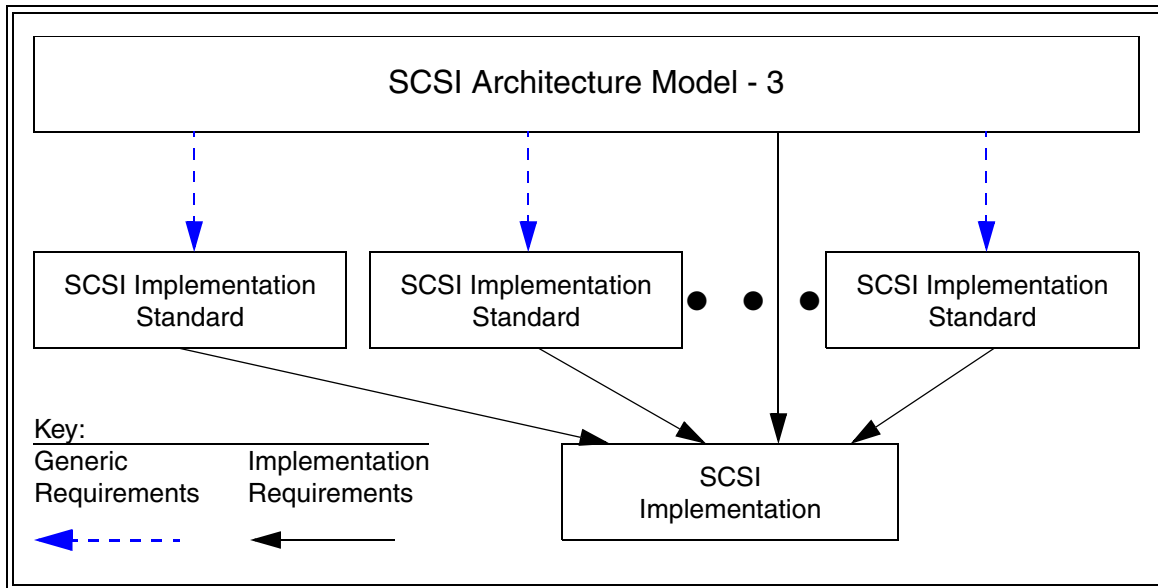


Figure 1 — Requirements precedence

As shown in figure 1, all SCSI implementation standards shall reflect the generic requirements defined herein. In addition, an implementation claiming SCSI compliance shall conform to the applicable implementation requirements defined in this standard and the appropriate SCSI implementation standards. In the event of a conflict between this document and other SCSI standards under the jurisdiction of technical committee T10, the requirements of this standard shall apply.

1.3 SCSI standards family

Figure 2 shows the relationship of this standard to the other standards and related projects in the SCSI family standards as of the publication of this standard.

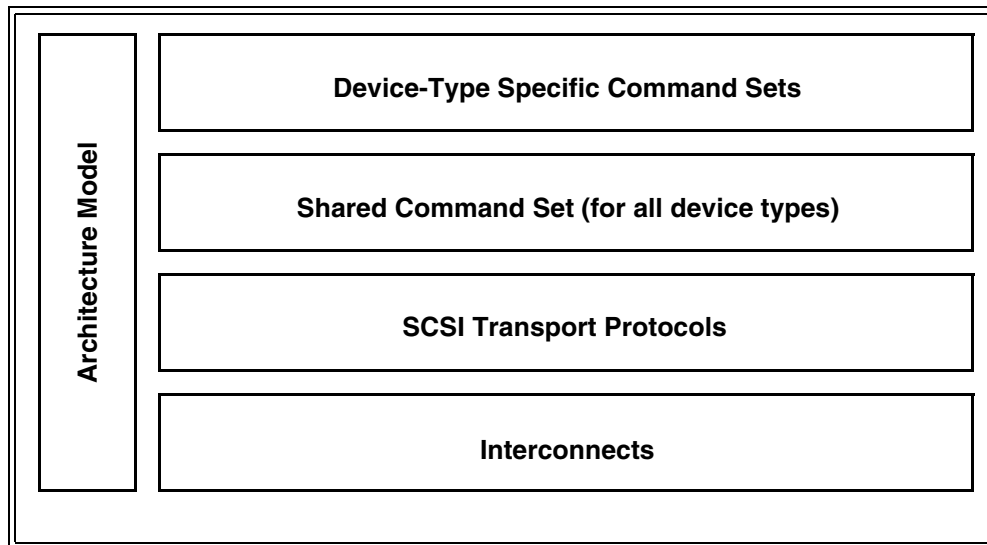


Figure 2 — SCSI document structure

The roadmap in figure 2 is intended to show the general applicability of the documents to one another. Figure 2 is not intended to imply a relationship such as a hierarchy, protocol stack, or system architecture.

The functional areas identified in figure 2 characterize the scope of standards within a group as follows:

Architecture Model: Defines the SCSI systems model, the functional partitioning of the SCSI standard set and requirements applicable to all SCSI implementations and implementation standards.

Device-Type Specific Command Sets: Implementation standards that define specific device types including a device model for each device type. These standards specify the required commands and behavior that is specific to a given device type and prescribe the requirements to be followed by a SCSI initiator device when sending commands to a SCSI target device having the specific device type. The commands and behaviors for a specific device type may include by reference commands and behaviors that are shared by all SCSI devices.

Shared Command Set: An implementation standard that defines a model for all SCSI device types. This standard specifies the required commands and behavior that is common to all SCSI devices, regardless of device type, and prescribes the requirements to be followed by a SCSI initiator device when sending commands to any SCSI target device.

SCSI Transport Protocols: Implementation standards that define the requirements for exchanging information so that different SCSI devices are capable of communicating.

Interconnects: Implementation standards that define the communications mechanism employed by the SCSI transport protocols. These standards may describe the electrical and signaling requirements essential for SCSI devices to interoperate over a given interconnect.

At the time this standard was generated, examples of the SCSI general structure included:

Interconnects:

Fibre Channel Arbitrated Loop - 2	FC-AL-2	[ISO/IEC 14165-122] [ANSI NCITS.332-1999] [ANSI NCITS.332-1999/AM1]
Fibre Channel Physical Interfaces	FC-PI	[ISO/IEC 14165-115] [ANSI INCITS.352-2002]
Fibre Channel Physical Interfaces - 2	FC-PI-2	[T11/1506-D]
Fibre Channel Framing and Signaling Interface	FC-FS	[ISO/IEC 14165-251] [ANSI INCITS.373-2003]
High Performance Serial Bus High Performance Serial Bus (supplement to ANSI/IEEE 1394-1995)		[ANSI IEEE 1394-1995] [ANSI IEEE 1394a-2000]
SCSI Parallel Interface - 2	SPI-2	[ISO/IEC 14776-112] [ANSI X3.302-1999]
SCSI Parallel Interface - 3	SPI-3	[ISO/IEC 14776-113] [ANSI NCITS.336-2000]
SCSI Parallel Interface - 4	SPI-4	[ISO/IEC 14776-114] [ANSI INCITS.362-2002]
SCSI Parallel Interface - 5	SPI-5	[ISO/IEC 14776-115] [ANSI INCITS.367:2003]
Serial Storage Architecture Physical Layer 1	SSA-PH	[ANSI X3.293-1996]
Serial Storage Architecture Physical Layer 2	SSA-PH-2	[ANSI NCITS.307-1998]
Serial Attached SCSI	SAS	[ISO/IEC 14776-150] [ANSI INCITS.376:2003]
Serial Attached SCSI - 1.1	SAS-1.1	[ISO/IEC 14776-151] [T10/1601-D]

SCSI Transport Protocols:

Automation/Drive Interface - Transport Protocol	ADT	[ISO/IEC 14776-191] [T10/1557-D]
Serial Storage Architecture Transport Layer 1	SSA-TL-1	[ANSI X3.295-1996]
Serial Storage Architecture Transport Layer 2	SSA-TL-2	[ANSI NCITS.308-1998]
SCSI-3 Fibre Channel Protocol	FCP	[ISO/IEC 14776-221] [ANSI X3.269-1996]
SCSI Fibre Channel Protocol - 2	FCP-2	[ISO/IEC 14776-222] [ANSI NCITS.350-2003]
SCSI Fibre Channel Protocol - 3	FCP-3	[ISO/IEC 14776-223] [T10/1560-D]
Serial Bus Protocol - 2	SBP-2	[ISO/IEC 14776-232] [ANSI NCITS.325-1999]
Serial Bus Protocol - 3	SBP-3	[ISO/IEC 14776-233] [T10/1467-D]
Serial Storage Architecture SCSI-3 Protocol	SSA-S3P	[ANSI NCITS.309-1998]
SCSI RDMA Protocol	SRP	[ISO/IEC 14776-241] [T10/1415-D]
SCSI RDMA Protocol - 2	SRP-2	[ISO/IEC 14776-242] [T10/1524-D]

Shared Command Sets:

SCSI-3 Primary Commands	SPC	[ANSI X3.301-1997]
SCSI Primary Commands - 2	SPC-2	[ISO/IEC 14776-452] [ANSI NCITS.351-2001]
SCSI Primary Commands - 3	SPC-3	[ISO/IEC 14776-453] [T10/1416-D]

Device-Type Specific Command Sets:

SCSI-3 Block Commands	SBC	[ISO/IEC 14776-321] [ANSI NCITS.306-1998]
SCSI Block Commands - 2	SBC-2	[ISO/IEC 14776-322] [T10/1417-D]
SCSI-3 Stream Commands	SSC	[ISO/IEC 14776-331] [ANSI NCITS.335-2000]
SCSI Stream Commands - 2	SSC-2	[ISO/IEC 14776-332] [ANSI INCITS.380-2003]
SCSI Stream Commands - 3	SSC-3	[ISO/IEC 14776-333] [T10/1611-D]
SCSI-3 Medium Changer Commands	SMC	[ISO/IEC 14776-351] [ANSI NCITS.314-1998]
SCSI Media Changer Commands - 2	SMC-2	[ISO/IEC 14776-352] [T10/1383-D]
SCSI-3 Multimedia Command Set	MMC	[ANSI X3.304-1997]
SCSI Multimedia Command Set - 2	MMC-2	[ISO/IEC 14776-362] [ANSI NCITS.333-2000]
SCSI Multimedia Command Set - 3	MMC-3	[ISO/IEC 14776-363] [ANSI INCITS.360-2002]
SCSI Multimedia Command Set - 4	MMC-4	[ISO/IEC 14776-364] [T10/1545-D]
SCSI Multimedia Command Set - 5	MMC-5	[ISO/IEC 14776-365] [T10/1xxx-D]
SCSI Controller Commands - 2	SCC-2	[ISO/IEC 14776-342] [ANSI NCITS.318-1998]
SCSI Reduced Block Commands	RBC	[ISO/IEC 14776-326] [ANSI NCITS.330-2000]
SCSI-3 Enclosure Services Commands	SES	[ISO/IEC 14776-371] [ANSI NCITS.305-1998]
SCSI Enclosure Services Commands - 2	SES-2	[ISO/IEC 14776-372] [T10/1559-D]
SCSI Specification for Optical Card Reader/Writer	OCRW	[ISO/IEC 14776-381]
Object-based Storage Devices Commands	OSD	[ISO/IEC 14776-391] [T10/1355-D]
SCSI Management Server Commands	MSC	[ISO/IEC 14776-511] [T10/1528-D]
Automation/Drive Interface - Commands	ADC	[ISO/IEC 14776-356] [T10/1558-D]

Architecture Model:

SCSI-3 Architecture Model	SAM	[ISO/IEC 14776-411] [ANSI X3.270-1996]
SCSI Architecture Model - 2	SAM-2	[ISO/IEC 14776-412] [ANSI INCITS.366-2003]
SCSI Architecture Model - 3	SAM-3	[ISO/IEC 14776-413] [T10/1561-D]

The term SCSI is used to refer to the family of standards described in this subclause.

2 Normative references

2.1 Normative references

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

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

2.2 Approved references

ISO/IEC 14776-452, SCSI Primary Commands - 2 (SPC-2) [ANSI NCITS.351-2001]

ISO/IEC 60027-2-am2 (1999-01), Letter symbols to be used in electrical technology - Part 2: Telecommunications and electronics (Amendment 2)

2.3 References under development

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

ISO/IEC 14776-453, SCSI Primary Commands - 3 (SPC-3) [T10/1416-D]

3 Definitions, symbols, abbreviations, and conventions

3.1 Definitions

3.1.1 ACA command: A command performed by a task with the ACA attribute (see 3.1.5, 4.11, and 8.6.5).

3.1.2 additional sense code: A combination of the ADDITIONAL SENSE CODE and ADDITIONAL SENSE CODE QUALIFIER fields in the sense data (see 3.1.109 and SPC-3).

3.1.3 application client: An object that is the source of commands.

3.1.4 argument: A datum provided as input to or output from a procedure call (see 3.1.77).

3.1.5 auto contingent allegiance (ACA): The task set condition established following the return of a CHECK CONDITION status when the NACA bit is set to one in the CONTROL byte. See 5.9.2.

3.1.6 background operation: An operation started by a command that continues processing after the task containing the command is no longer in the task set. See 5.5.

3.1.7 basic task management model: The task management model in which only one task attribute (see 8.6) is supported and the task management features that the application client may select are limited. See 8.3.3.

3.1.8 blocked task state: When in this state a task is prevented from completing due to an ACA condition.

3.1.9 blocking boundary: A task set boundary denoting a set of conditions that inhibit tasks outside the boundary from entering the enabled task state.

3.1.10 byte: An 8-bit construct.

3.1.11 client-server: A relationship established between a pair of distributed entities where one (the client) requests the other (the server) to perform some operation or unit of work on the client's behalf.

3.1.12 client: An entity that requests a service from a server. This standard defines one client, the application client.

3.1.13 code value: A defined numeric value, possibly a member of a series of defined numeric values, representing an identified and described instance or condition. Code values are defined to be used in a specific field (see 3.1.36), in a procedure call input argument (see 3.6.2), in a procedure call output argument, or in a procedure call result.

3.1.14 command: A request describing a unit of work to be performed by a device server.

3.1.15 command descriptor block (CDB): A structure used to communicate a command from an application client to a device server. A CDB may have a fixed length of up to 16 bytes or a variable length of between 12 and 260 bytes.

3.1.16 command standard: A SCSI standard that defines the model, commands, and parameter data for a device type (e.g., SPC-3, SBC-2, SSC-2, SMC-2, MMC-3, or SES-2). See clause 1.

3.1.17 completed command: A command that has ended by returning a status and service response of TASK COMPLETE or LINKED COMMAND COMPLETE.

3.1.18 completed task: A task that has ended by returning a status and service response of TASK COMPLETE. The actual events comprising the TASK COMPLETE response are SCSI transport protocol specific.

3.1.19 confirmation: A response returned to an application client or device server that signals the completion of a service request.

3.1.20 confirmed SCSI transport protocol service: A service available at the SCSI transport protocol service interface that includes a confirmation of completion.

3.1.21 current task: A task that has a data transfer SCSI transport protocol service request in progress (see 5.4.3) or is in the process of sending command status. Each SCSI transport protocol standard should define the SCSI transport protocol specific conditions under which a task is considered a current task.

3.1.22 deferred error: An error generated by a background operation (see 3.1.6).

3.1.23 dependent logical unit: A logical unit that is addressed via some other logical unit(s) in a hierarchical logical unit structure (see 3.1.40), also a logical unit that is at a higher numbered level in the hierarchy than the referenced logical unit (see 4.14).

3.1.24 device identifier: A term used by previous versions of this standard (see Annex B).

3.1.25 device model: The description of a type of SCSI target device (e.g., block, stream).

3.1.26 device server: An object within a logical unit that processes SCSI tasks according to the requirements for task management described in clause 8.

3.1.27 device service request: A request submitted by an application client conveying a command to a device server.

3.1.28 device service response: The response returned to an application client by a device server on completion of a command.

3.1.29 domain: An I/O system consisting of a set of SCSI devices that interact with one another by means of a service delivery subsystem.

3.1.30 dormant task state: When in this state a task is prevented from entering the enabled task state (see 3.1.31) due to the presence of certain other tasks in the task set.

3.1.31 enabled task state: When in this state a task may complete at any time or is waiting to receive the next command in a series of linked commands.

3.1.32 faulted initiator port: The SCSI initiator port to which a CHECK CONDITION status was returned that resulted in the establishment of an ACA. The faulted initiator port condition is cleared when the ACA condition is cleared.

3.1.33 faulted task set: A task set that contains a faulting task. The faulted task set condition is cleared when the ACA condition resulting from the CHECK CONDITION status is cleared.

3.1.34 faulting command: A command that completed with a status of CHECK CONDITION that resulted in the establishment of an ACA.

3.1.35 faulting task: A task that has completed with a status of CHECK CONDITION that resulted in the establishment of an ACA.

3.1.36 field: A group of one or more contiguous bits, part of a larger structure (e.g., a CDB (see 3.1.15) or sense data (see 3.1.109)).

3.1.37 full task management model: The task management model in which the SIMPLE task attribute (see 8.6.1) is supported and all task management features may be selected by the application client. See 8.3.2.

3.1.38 function complete: A logical unit response indicating that a task management function has finished. The events comprising this response are SCSI transport protocol specific.

3.1.39 hard reset: A response to a power on (see 6.3.1) or reset (see 6.3.2) condition in which the SCSI device performs the operations described in 6.3.2.

3.1.40 hierarchical logical unit: An inverted tree structure for forming and parsing logical unit numbers (see 3.1.64) containing up to four addressable levels (see 4.14).

3.1.41 I_T nexus: A nexus between a SCSI initiator port and a SCSI target port (see 4.12).

3.1.42 I_T nexus loss: A condition resulting from a hard reset condition (see 6.3.2) in SCSI initiator devices or from delivery of an I_T nexus loss notification indication (see 6.4) in all SCSI devices in which the SCSI device performs the operations described in 6.3.4.

3.1.43 I_T nexus loss notification: An indication (see 6.4) from the SCSI transport protocol to the SCSI application layer that an I_T nexus no longer exists.

3.1.44 I_T_L nexus: A nexus between a SCSI initiator port, a SCSI target port, and a logical unit (see 4.12).

3.1.45 I_T_L_Q nexus: A nexus between a SCSI initiator port, a SCSI target port, a logical unit, and a task (see 4.12).

3.1.46 I_T_L_Q nexus transaction: The information transferred between SCSI ports in a single data structure with defined boundaries (e.g., an information unit).

3.1.47 I_T_L_x nexus: Either an I_T_L nexus or an I_T_L_Q nexus (see 4.12).

3.1.48 I/O operation: An operation defined by an unlinked command, a series of linked commands or a task management function.

3.1.49 implementation specific: A requirement or feature that is defined in a SCSI standard but whose implementation may be specified by the system integrator or vendor.

3.1.50 initiator: A term used by previous versions of this standard (see Annex B).

3.1.51 initiator device name: A SCSI device name of a SCSI initiator device (see 4.7.1).

3.1.52 initiator identifier: A term used by previous versions of this standard (see Annex B).

3.1.53 initiator port identifier: A value by which a SCSI initiator port is referenced within a domain. See 4.7.1.

3.1.54 initiator port name: A SCSI port name (see 3.1.98) of a SCSI initiator port or of a SCSI target/initiator port when operating as a SCSI initiator port. See 4.7.1.

3.1.55 interconnect subsystem: One or more interconnects that appear as a single path for the transfer of information between SCSI devices in a domain.

3.1.56 in transit: Information that has been sent to a remote entity but not yet received.

3.1.57 implicit head of queue: An optional processing model for specified commands wherein the specified commands may be treated as if they had been received with a HEAD OF QUEUE task attribute. See 8.2.

3.1.58 layer: A subdivision of the architecture constituted by SCSI initiator device and SCSI target device elements at the same level relative to the interconnect.

3.1.59 linked command: One in a series of commands processed by a single task that collectively make up a discrete I/O operation. In such a series, each command is represented by the same I_T_L_Q nexus, and all, except the last, have the LINK bit in the CDB CONTROL byte set to one.

3.1.60 logical unit: A SCSI target device object, containing a device server and task manager, that implements a device model and manages tasks to process commands sent by an application client. See 4.8.

3.1.61 logical unit reset: A logical unit response to a logical unit reset event in which the logical unit performs the operations described in 6.3.3.

3.1.62 logical unit reset event: An event that triggers a logical unit reset (see 3.1.61). Logical unit reset events include processing the LOGICAL UNIT RESET task management function (see 7.6) and hard reset (see 6.3.2).

3.1.63 logical unit inventory: The list of the logical unit numbers reported by a REPORT LUNS command (see SPC-3).

3.1.64 logical unit number (LUN): A 64-bit identifier for a logical unit. See 4.9.

3.1.65 media information: Information stored within a SCSI device that is non-volatile (retained through a power cycle) and accessible to a SCSI initiator device through the processing of commands.

3.1.66 name: A label of an object that is unique within a specified context and should never change (e.g., the term name and world wide identifier (WWID) may be interchangeable).

3.1.67 nexus: A relationship between two SCSI devices, and the SCSI initiator port and SCSI target port objects within those SCSI devices. See 4.12.

3.1.68 non-faulted initiator port: A SCSI initiator port that is not a faulted initiator port (see 3.1.32).

3.1.69 object: A container that encapsulates data types, services, or other objects that are related in some way.

3.1.70 peer entities: Entities within the same layer.

3.1.71 pending command: From the point of view of the application client, the description of command between the time that the application client calls the **Send SCSI Command** SCSI transport protocol service and the time one of the SCSI target device responses described in 5.5 is received.

3.1.72 port: Synonymous with SCSI port (see 3.1.96).

3.1.73 power cycle: Power being removed from and later applied to a SCSI device.

3.1.74 power on: A condition resulting from a power on event (see 3.1.75).

3.1.75 power on event: Power being applied to a SCSI device, triggering a power on condition (see 3.1.74) in the SCSI device.

3.1.76 procedure: An operation that is invoked through an external calling interface.

3.1.77 procedure call: The model used by this standard for the interfaces involving both the SAL (see 3.1.88) and STPL (see 3.1.103), having the appearance of a programming language function call.

3.1.78 protocol: A specification and/or implementation of the requirements governing the content and exchange of information passed between distributed entities through the service delivery subsystem.

3.1.79 queue: The arrangement of tasks within a task set (see 3.1.131), usually according to the temporal order in which they were created.

3.1.80 receiver: A client or server that is the recipient of a service delivery transaction.

3.1.81 reference model: A standard model used to specify system requirements in an implementation-independent manner.

3.1.82 relative port identifier: An identifier for a SCSI port that is unique within a SCSI device. See 4.7.5.

3.1.83 request: A transaction invoking a service.

3.1.84 request-response transaction: An interaction between a pair of distributed, cooperating entities, consisting of a request for service submitted to an entity followed by a response conveying the result.

3.1.85 request-confirmation transaction: An interaction between a pair of cooperating entities, consisting of a request for service submitted to an entity followed by a response from the entity confirming request completion.

3.1.86 reset event: A SCSI transport protocol specific event that triggers a hard reset (see 6.3.2).

3.1.87 response: A transaction conveying the result of a request.

3.1.88 SCSI application layer (SAL): The protocols and procedures that implement or issue commands and task management functions by using services provided by a SCSI transport protocol layer.

3.1.89 SCSI device: A device that contains one or more SCSI ports that are connected to a service delivery subsystem and supports a SCSI application protocol.

3.1.90 SCSI device identifier: Synonymous with SCSI port identifier (see 3.1.97).

3.1.91 SCSI device name: A name (see 3.1.66) of a SCSI device that is world wide unique within the SCSI transport protocol of a SCSI domain in which the SCSI device has SCSI ports (see 4.7.7). The SCSI device name may be made available to other SCSI devices or SCSI ports in SCSI transport protocol specific ways.

3.1.92 SCSI I/O system: An I/O system, consisting of two or more SCSI devices, a SCSI interconnect and a SCSI transport protocol that collectively interact to perform SCSI I/O operations.

3.1.93 SCSI identifier: A term used by previous versions of this standard (see Annex B).

3.1.94 SCSI initiator device: A SCSI device containing application clients and SCSI initiator ports that originates device service and task management requests to be processed by a SCSI target device and receives device service and task management responses from SCSI target devices. When used this term refers to SCSI initiator devices or SCSI target/initiator devices that are using the SCSI target/initiator port as a SCSI initiator port.

3.1.95 SCSI initiator port: A SCSI initiator device object that acts as the connection between application clients and the service delivery subsystem through which requests and confirmations are routed. In all cases when this term is used it refers to an initiator port or a SCSI target/initiator port operating as a SCSI initiator port.

3.1.96 SCSI port: A SCSI device resident object that connects the application client, device server or task manager to the service delivery subsystem through which requests and responses are routed. SCSI port is synonymous with port. A SCSI port is either a SCSI initiator port (see 3.1.95) or a SCSI target port (see 3.1.100).

3.1.97 SCSI port identifier: A value by which a SCSI port is referenced within a domain. The SCSI port identifier is either an initiator port identifier (see 3.1.53) or a target port identifier (see 3.1.121).

3.1.98 SCSI port name: A name (see 3.1.66) of a SCSI port that is world wide unique within the SCSI transport protocol of the SCSI domain of that SCSI port (see 4.7.8). The name may be made available to other SCSI devices or SCSI ports in that SCSI domain in SCSI transport protocol specific ways.

3.1.99 SCSI target device: A SCSI device containing logical units and SCSI target ports that receives device service and task management requests for processing and sends device service and task management responses to SCSI initiator devices. When used this term refers to SCSI target devices or SCSI target/initiator devices that are using the SCSI target/initiator port as a SCSI target port.

3.1.100 SCSI target port: A SCSI target device object that contains a task router and acts as the connection between device servers and task managers and the service delivery subsystem through which indications and responses are routed. When this term is used it refers to a SCSI target port or a SCSI target/initiator port operating as a SCSI target port.

3.1.101 SCSI target/initiator device: A SCSI device that has all the characteristics of a SCSI target device and a SCSI initiator device.

3.1.102 SCSI target/initiator port: A SCSI device resident object that has all the characteristics of a SCSI target port and a SCSI initiator port.

3.1.103 SCSI transport protocol layer (STPL): The protocol and services used by a SCSI application layer to transport data representing a SCSI application protocol transaction.

3.1.104 SCSI transport protocol service confirmation: A procedure call from the STPL notifying the SAL that a SCSI transport protocol service request has completed.

3.1.105 SCSI transport protocol service indication: A procedure call from the STPL notifying the SAL that a SCSI transport protocol transaction has occurred.

3.1.106 SCSI transport protocol service request: A procedure call to the STPL to begin a SCSI transport protocol service transaction.

3.1.107 SCSI transport protocol service response: A procedure call to the STPL containing a reply from the SAL in response to a SCSI transport protocol service indication.

3.1.108 sender: A client or server that originates a service delivery transaction.

3.1.109 sense data: Data returned to an application client in the same I_T_L_Q nexus transaction (see 3.1.46) as a CHECK CONDITION status (see 5.9.6). Fields in the sense data are referenced by name in this standard. See SPC-3 for a complete sense data format definition. Sense data may also be retrieved using the REQUEST SENSE command (see SPC-3).

3.1.110 sense key: A field in the sense data (see 3.1.109 and SPC-3).

- 3.1.111 server:** An entity that performs a service on behalf of a client.
- 3.1.112 service:** Any operation or function performed by a SCSI object that is invoked by other SCSI objects.
- 3.1.113 service delivery failure:** Any non-recoverable error causing the corruption or loss of one or more service delivery transactions while in transit.
- 3.1.114 service delivery subsystem:** That part of a SCSI I/O system that transmits service requests to a logical unit or SCSI target device and returns logical unit or SCSI target device responses to a SCSI initiator device.
- 3.1.115 service delivery transaction:** A request or response sent through the service delivery subsystem.
- 3.1.116 signal:** (n) A detectable asynchronous event possibly accompanied by descriptive data and parameters.
(v) The act of generating such an event.
- 3.1.117 standard INQUIRY data:** Data returned to an application client as a result of an INQUIRY command. Fields in the standard INQUIRY data are referenced by name in this standard and SPC-3 contains a complete definition of the standard INQUIRY data format.
- 3.1.118 target:** A term used by previous versions of this standard (see Annex B).
- 3.1.119 target device name:** A SCSI device name (see 3.1.91) of a SCSI target device. See 4.7.2.
- 3.1.120 target identifier:** A term used by previous versions of this standard (see Annex B).
- 3.1.121 target port identifier:** A value by which a SCSI target port is referenced within a domain. See 4.7.2.
- 3.1.122 target port name:** A SCSI port name of a SCSI target port or of a SCSI target/initiator port when operating as a SCSI target port (see 4.7.2).
- 3.1.123 target/initiator device name:** A SCSI device name (see 3.1.91) of a SCSI target/initiator device. See 4.7.3.
- 3.1.124 task:** An object within the logical unit representing the work associated with a command or a group of linked commands.
- 3.1.125 task tag:** An object containing up to 64 bits that is a component of an I_T_L_Q nexus. See 4.11.
- 3.1.126 task management function:** A task manager service capable of being requested by an application client to affect the processing of one or more tasks.
- 3.1.127 task management request:** A request submitted by an application client, invoking a task management function to be processed by a task manager.
- 3.1.128 task management response:** The response returned to an application client by a task manager on completion of a task management request.
- 3.1.129 task manager:** A server within a logical unit that controls the sequencing of one or more tasks and processes task management functions.
- 3.1.130 task router:** An object in a SCSI target port that routes commands and task management functions between the service delivery subsystem (see 3.1.114) and the appropriate logical unit's task manager (see 3.1.129).

3.1.131 task set: A group of tasks within a logical unit, whose interaction is dependent on the task management (e.g., queuing) and ACA requirements. See 4.8.

3.1.132 third-party command: A command that requires a logical unit within a SCSI target device to assume the SCSI initiator device role and send command(s) to another SCSI target device.

3.1.133 transaction: A cooperative interaction between two entities, involving the exchange of information or the processing of some request by one entity on behalf of the other.

3.1.134 unconfirmed SCSI transport protocol service: A service available at the SCSI transport protocol service interface that does not result in a completion confirmation.

3.1.135 unlinked command: A command having the LINK bit set to zero in the CDB CONTROL byte.

3.1.136 well known logical unit: A logical unit that only performs specific functions. See 4.9.9. Well known logical units allow an application client to issue requests to receive and manage specific information usually relating to a SCSI target device.

3.1.137 well known logical unit number (W-LUN): The logical unit number that identifies a well known logical unit.

3.2 Acronyms

ACA	Auto Contingent Allegiance (see 3.1.5)
CDB	Command Descriptor Block (see 3.1.15)
LUN	Logical Unit Number (see 3.1.64)
MMC-2	SCSI Multi-Media Commands -2 (see 1.3)
n/a	Not Applicable
RAID	Redundant Array of Independent Disks
SAL	SCSI application layer (see 3.1.88)
SBC	SCSI-3 Block Commands (see 1.3)
SCSI	The architecture defined by the family of standards described in 1.3
SPI-5	SCSI Parallel Interface -5 (see 1.3)
SPC-3	SCSI Primary Commands -3 (see 1.3)
STPL	SCSI transport protocol layer (see 3.1.103)
SSC	SCSI-3 Stream Commands (see 1.3)
VPD	Vital Product Data (see SPC-3)
W-LUN	Well known logical unit number (see 3.1.137)

3.3 Keywords

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

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

3.3.3 may: A keyword that indicates flexibility of choice with no implied preference (synonymous with "may or may not").

3.3.4 may not: A keyword that indicates flexibility of choice with no implied preference (synonymous with "may or may not").

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

3.3.6 option, optional: Keywords that describe features that are not required to be implemented by this standard. However, if any optional feature defined by this standard is implemented, then it shall be implemented as defined in this standard.

3.3.7 SCSI transport protocol specific: Implementation of the referenced item is defined by a SCSI transport protocol standard (see 1.3).

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

3.3.9 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

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

3.3.11 vendor specific: Specification of the referenced item is determined by the SCSI device vendor.

3.4 Editorial conventions

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

Upper case is used when referring to the name of a numeric value defined in this specification or a formal attribute possessed by an entity. When necessary for clarity, names of objects, procedure calls, arguments or discrete states are capitalized or set in bold type. Names of fields are identified using small capital letters (e.g., NACA bit).

Names of procedure calls are identified by a name in bold type, such as **Execute Command** (see clause 5). Names of arguments are denoted by capitalizing each word in the name. For instance, Sense Data is the name of an argument in the **Execute Command** procedure call.

Quantities having a defined numeric value are identified by large capital letters. CHECK CONDITION, for example, refers to the numeric quantity defined in table 21 (see 5.3.1). Quantities having a discrete but unspecified value are identified using small capital letters. As an example, TASK COMPLETE, indicates a quantity returned by the **Execute Command** procedure call (see clause 5). Such quantities are associated with an event or indication whose observable behavior or value is specific to a given implementation standard.

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

If a conflict arises between text, tables, or figures, the order of precedence to resolve the conflicts is text; then tables; and finally figures. Not all tables or figures are fully described in the text. Tables show data format and values.

Notes do not constitute any requirements for implementors.

3.5 Numeric conventions

Digits 0 through 9 in the text of this standard that are not immediately followed by lower-case "b" or "h" are decimal values. Digits 0 and 1 immediately followed by lower case "b" are binary values. Digits 0 through 9 and the upper case letters "A" through "F" immediately followed by lower-case "h" are hexadecimal values.

Large numbers are separated by spaces (e.g., 12 345, not 12,345).

3.6 Notation conventions

3.6.1 Hierarchy diagram conventions

Hierarchy diagrams show how objects are related to each other. The hierarchy diagram of figure 3, for example, shows the relationships among the objects comprising an object called Book. For this example, a book object is defined as containing a table of contents object, an optional preface object, one or more chapter objects, and an optional index object. Further contents definitions are provided for the preface and chapter objects. A preface object contains zero or more figure objects, one outline object, and an introductory text object. A chapter object contains one or more section objects and zero or more figure objects.

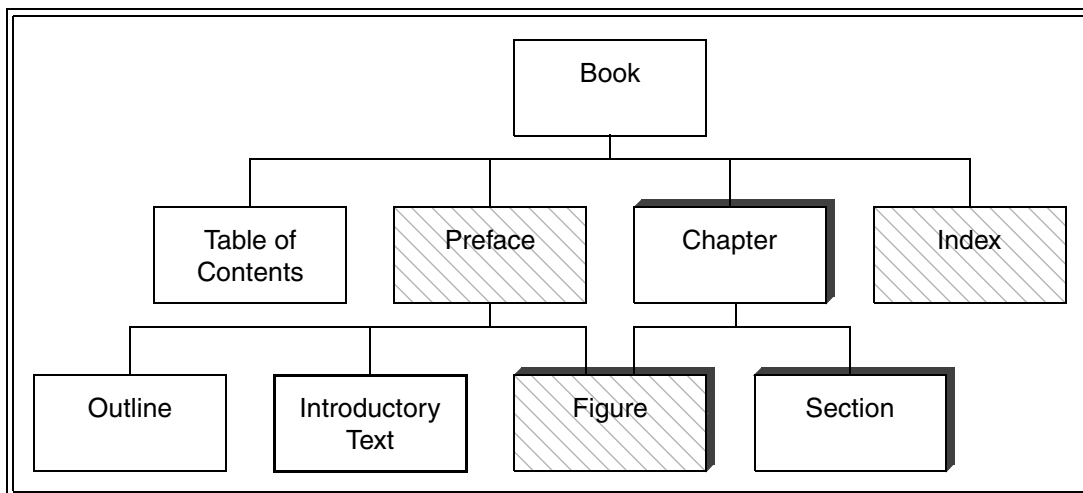


Figure 3 — Example hierarchy diagram

In the corresponding hierarchy diagram, labeled boxes denote the above objects. The composition and relation of one object to others is shown by the connecting lines. In this case, the connecting lines indicate the relationship between the book object and its constituent table of contents, preface, chapter and index objects. Similarly, connecting lines show that a chapter object contains section and figure objects. Note that the figure object also may be a component of the preface object.

In the hierarchy diagram, objects that are required to have one and only one instance are shown as simple boxes, as is the case for the book and table of contents objects. The hierarchy diagram shows multiple instances of an object by the presence of a shadow, as is the case for the chapter, figure and section objects. Objects that are optional are indicated by light diagonal lines, as is the case for the preface, figure and index objects. An object that may not have any instances, have only one instance, or have multiple instances is shown with both diagonal lines and a shadow, as is the case for the figure object. The instance indications shown in a hierarchy diagram are approximate; detailed requirements appear in the accompanying text.

3.6.2 Notation for procedure calls

In this standard, the model for functional interfaces between entities is a procedure call (see 3.1.77). Such interfaces are specified using the following notation:

[Result =] Procedure Name (IN ([input-1] [,input-2] ...), OUT ([output-1] [,output-2] ...))

Where:

Result: A single value representing the outcome of the procedure call.

Procedure Name: A descriptive name for the function modeled by the procedure call.

Input-1, Input-2, ...: A comma-separated list of names identifying caller-supplied input arguments.

Output-1, Output-2, ...: A comma-separated list of names identifying output arguments to be returned by the procedure call.

"[...]": Brackets enclosing optional or conditional arguments.

This notation allows arguments to be specified as inputs and outputs. The following is an example of a procedure call specification:

Found = Search (IN (Pattern, Item List), OUT ([Item Found]))

Where:

Found = Flag

Flag, if set to one, indicates that a matching item was located.

Input Arguments:

Pattern = ... /* Definition of **Pattern** argument */
Argument containing the search pattern.

Item List = Item<NN> /* Definition of **Item List** as an array of NN **Item** arguments*/
Contains the items to be searched for a match.

Output Arguments:

Item Found = Item ... /* Item located by the search procedure call */
This argument is only returned if the search succeeds.

3.6.3 Notation for state diagrams

All state diagrams use the notation shown in figure 4.

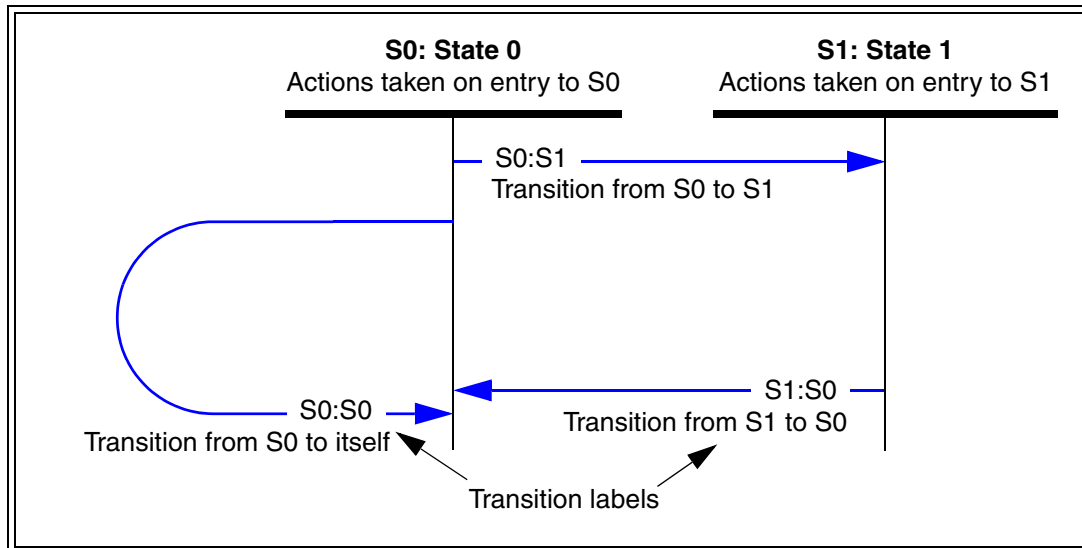


Figure 4 — Example state diagram

The state diagram is followed by a list of the state transitions, using the transition labels. Each transition is described in the list with particular attention to the conditions that cause the transition to occur and special conditions related to the transition. Using figure 4 as an example, the transition list might read as follows:

Transition S0:S1: This transition occurs when state S0 is exited and state S1 is entered.

Transition S1:S0: This transition occurs when state S1 is exited and state S0 is entered.

Transition S0:S0: This transition occurs when state S0 transitions to itself. The reason for a transition from S0 to itself is to specify that the actions taken whenever state S0 is entered are repeated every time the transition occurs.

A system specified in this manner has the following properties:

- a) Time elapses only within discrete states;
- b) State transitions are logically instantaneous; and
- c) Every time a state is entered, the actions of that state are started. Note that this means that a transition that points back to the same state restarts the actions from the beginning.

4 SCSI architecture model

4.1 Introduction

The purpose of the SCSI architecture model is to:

- a) Provide a basis for the coordination of SCSI standards development that allows each standard to be placed into perspective within the overall SCSI Architecture model;
- b) Identify areas for developing standards and provide a common reference for maintaining consistency among related standards so that independent implementors may work productively and independently on the development of standards within each functional area; and
- c) Provide the foundation for application compatibility across all SCSI interconnect and SCSI transport protocol environments by specifying generic requirements that apply uniformly to all implementation standards within each functional area.

The development of this standard is assisted by the use of an abstract model. To specify the external behavior of a SCSI system, elements in a system are replaced by functionally equivalent components within this model. Only externally observable behavior is retained as the standard of behavior. The description of internal behavior in this standard is provided only to support the definition of the observable aspects of the model. Those aspects are limited to the generic properties and characteristics needed for host applications to interoperate with SCSI devices in any SCSI interconnect and SCSI transport protocol environment. The model does not address other requirements that may be essential to some I/O system implementations (e.g., the mapping from SCSI device addresses to network addresses, the procedure for discovering SCSI devices on a network, and the definition of network authentication policies for SCSI initiator devices or SCSI target devices). These considerations are outside the scope of this standard.

The set of SCSI standards specifies the interfaces, functions, and operations necessary to ensure interoperability between conforming SCSI implementations. This standard is a functional description. Conforming implementations may employ any design technique that does not violate interoperability.

The SCSI architecture model is described in terms of objects (see 3.1.69), protocol layers, and service interfaces between objects. As used in this standard, objects are abstractions, encapsulating a set of related functions, data types, and other objects. Certain objects are defined by SCSI (e.g., an interconnect), while others are needed to understand the functioning of SCSI but have implementation definitions outside the scope of SCSI (e.g., a task). These objects exhibit well-defined and observable behaviors, but they do not exist as separate physical elements. An object may be a single numeric parameter (e.g., a task tag) or a complex entity that performs a set of operations or services on behalf of another object.

Service interfaces are defined between distributed objects and protocol layers. The template for a distributed service interface is the client-server model described in 4.2. The structure of a SCSI I/O system is specified in 4.4 by defining the relationship among objects. The set of distributed services to be provided are specified in clause 5 and clause 7.

Requirements that apply to each SCSI transport protocol standard are specified in the SCSI transport protocol service model described in 5.4 and 7.8. The model describes required behavior in terms of layers, objects within layers and SCSI transport protocol service transactions between layers.

4.2 The SCSI distributed service model

Service interfaces between distributed objects are represented by the client-server model shown in figure 5. Dashed horizontal lines with arrowheads denote a single request-response transaction as it appears to the client and server. The solid lines with arrowheads indicate the actual transaction path through the service delivery subsystem. In such a model, each client or server is a single thread of processing that runs concurrently with all other clients or servers.

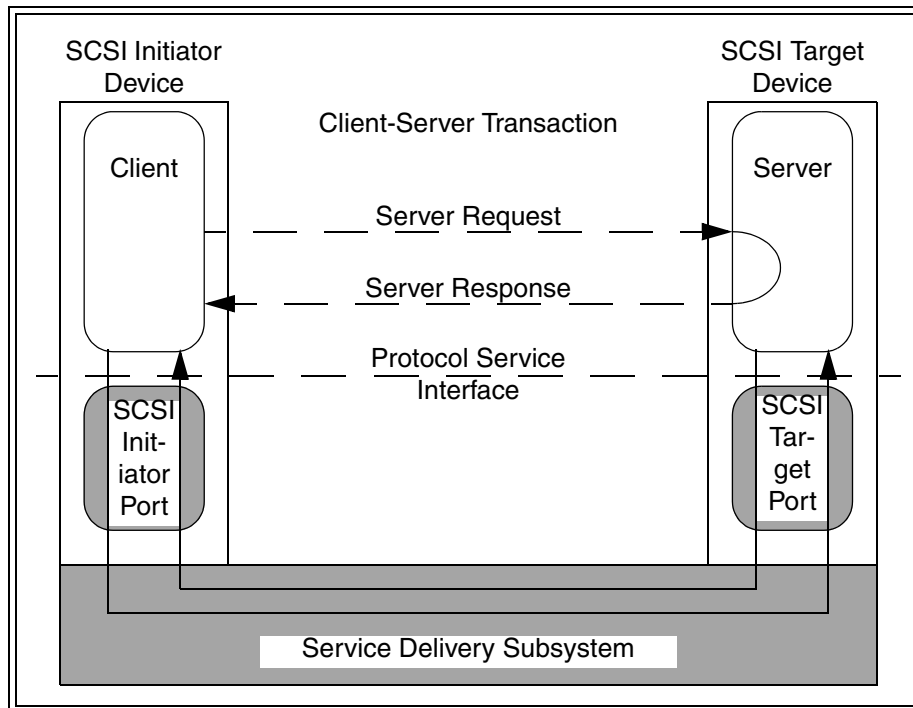


Figure 5 — Client-Server model

A client-server transaction is represented as a procedure call with inputs supplied by the caller (i.e., the client). The procedure call is processed by the server and returns outputs and a procedure call status. A client directs requests to a remote server via the SCSI initiator port and service delivery subsystem and receives a completion response or a failure notification. The request identifies the server and the service to be performed and includes the input data. The response conveys the output data and request status. The function of the service delivery subsystem is to transport error-free copies of the requests and responses between sender and receiver. A failure notification indicates that a condition has been detected (e.g., a reset or service delivery failure) that precludes request completion.

As seen by the client, a request becomes pending when it is passed to the SCSI initiator port for transmission. The request is complete when the server response is received or when a failure notification is sent. As seen by the server, the request becomes pending upon receipt and completes when the response is passed to the SCSI target port for return to the client. As a result there may be a time skew between the server and client's perception of request status and server state. All references to a pending command or task management function in this standard are from the application client's point of view (see 5.5 and 5.6).

Client-server relationships are not symmetrical. A client may only originate requests for service. A server may only respond to such requests. The client requests an operation provided by a server located in another SCSI device and waits for completion, which includes transmission of the request and response to/from the remote server. From the client's point of view, the behavior of a service requested from another SCSI device is indistinguishable from a request processed in the same SCSI device. In this model, confirmation of successful request or response delivery

by the sender is not required. The model assumes that delivery failures are detected by the SCSI initiator port or within the service delivery subsystem.

4.3 The SCSI client-server model

As shown in figure 6, each SCSI target device contains one or more logical units and provides services performed by device servers and task management functions performed by task managers. A logical unit is an object that implements one of the device functional models described in the SCSI command standards and processes commands (e.g., reading from or writing to the media). Each pending command or series of linked commands defines a unit of work to be performed by the logical unit. Each unit of work is represented within the SCSI target device by a task that may be externally referenced and controlled through requests issued to the task manager.

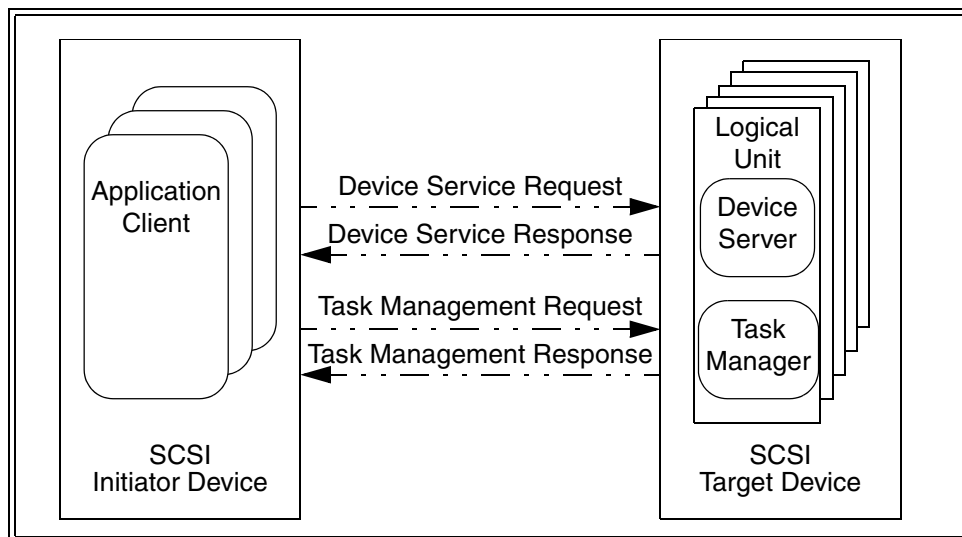


Figure 6 — SCSI client-server model

All requests originate from application clients residing within a SCSI initiator device. An application client is independent of the interconnect and SCSI transport protocol (e.g., an application client may correspond to the device driver and any other code within the operating system that is capable of managing I/O requests without requiring knowledge of the interconnect or SCSI transport protocol). An application client creates one or more application client tasks each of which issues a single command, series of linked commands, or task management function. Application client tasks are part of their parent application client. An application client task ceases to exist once the command, series of linked commands, or task management function ends. There is one application client task for each pending command, series of linked commands, or task management request.

As described in 4.2, each request takes the form of a procedure call with arguments and a status to be returned. An application client may request processing of a command through a request directed to the device server within a logical unit. Each device service request contains a CDB defining the operation to be performed along with a list of command specific inputs and other parameters specifying how the command is to be processed. If supported by a logical unit, a series of linked commands may be used to define an extended I/O operation.

A task is an object within the logical unit representing the work associated with a command or series of linked commands. A new command or the first in a series of linked commands causes the creation of a task. The task persists until a command completion response is sent or until the task is ended by a task management function or exception condition. For an example of the processing for a single command see 5.8.1. For an example of linked command processing see 5.8.2.

An application client may request processing of a task management function through a request directed to the task manager within the logical unit. The interactions between the task manager and application client when a task management request is processed are shown in 7.9.

4.4 The SCSI structural model

The SCSI structural model represents a view of the elements comprising a SCSI I/O system as seen by the application clients interacting with the system. As shown in figure 7, the fundamental object is the SCSI domain that represents an I/O system. A domain is made up of SCSI devices and a service delivery subsystem that transports commands, data, task management functions, and related information. A SCSI device contains clients or servers or both and the infrastructure to support them.

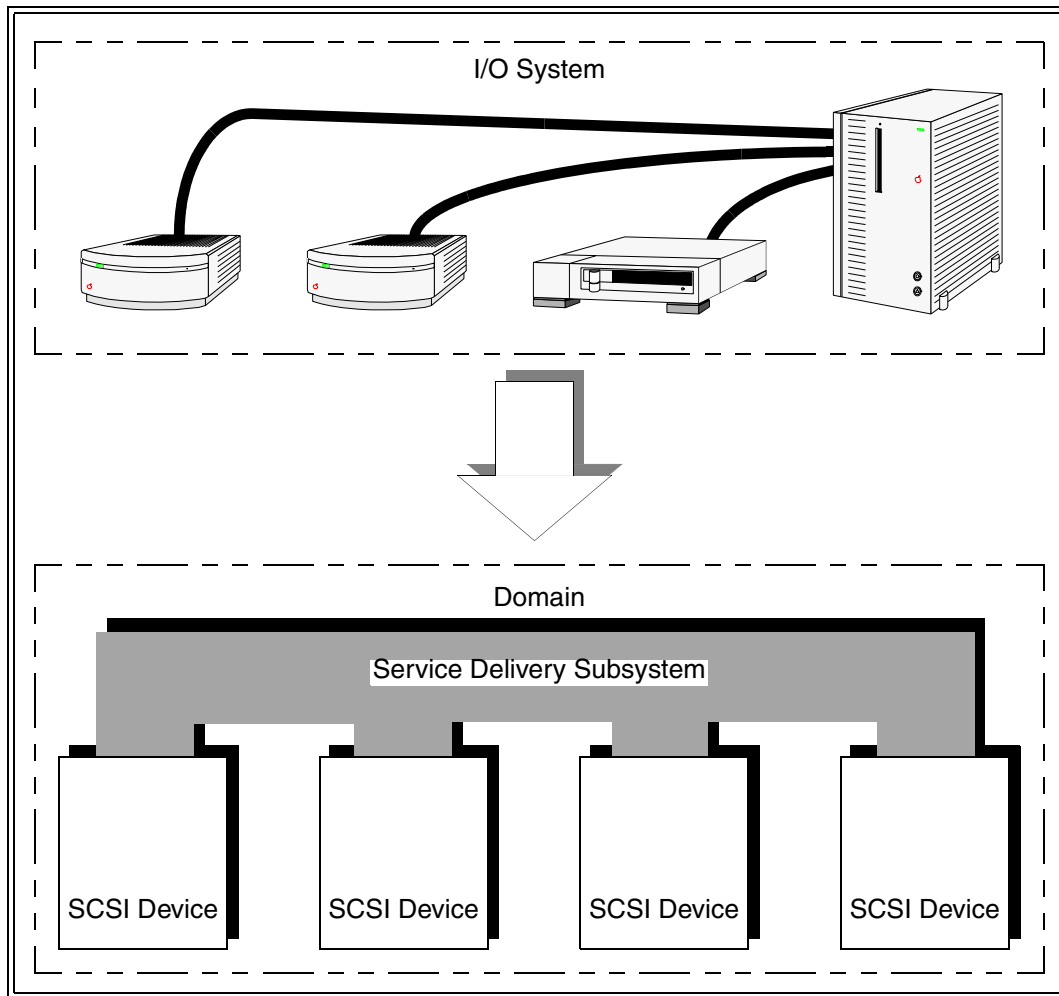


Figure 7 — SCSI I/O system and domain model

Figure 8 shows the main functional components of the SCSI domain. This standard defines these components in greater detail.

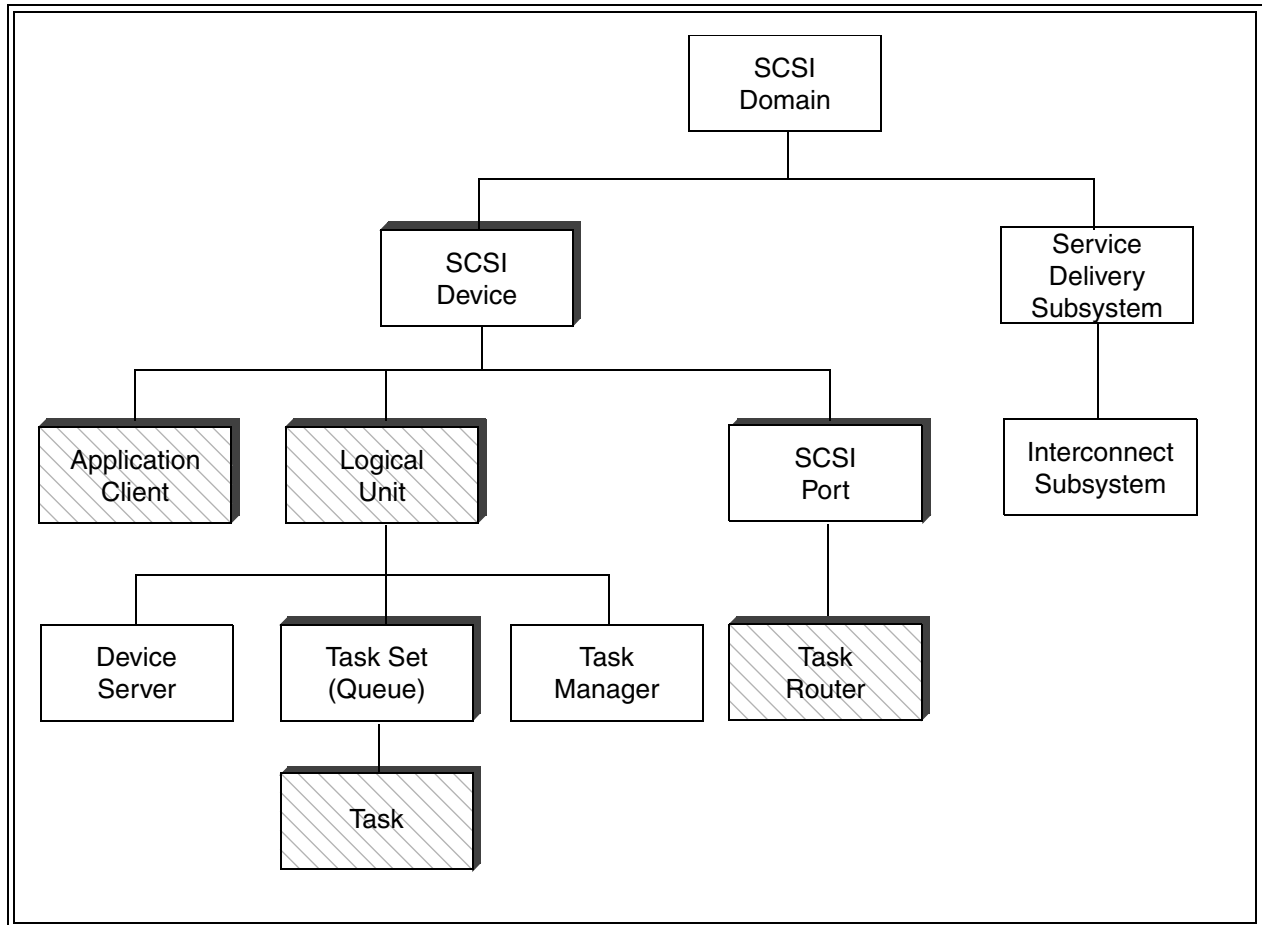


Figure 8 — Overall SCSI domain model

4.5 SCSI domain

A SCSI domain is composed of at least one SCSI device, at least one SCSI target port and at least one SCSI initiator port interconnected by a service delivery subsystem (see figure 9).




Figure 9 — SCSI domain model

A SCSI device is an object that originates or processes commands (see 4.7).

When a SCSI device originates a command it is called a SCSI initiator device. The commands are transmitted through SCSI initiator ports or SCSI target/initiator ports.

A SCSI device containing logical units that process commands is called a SCSI target device. It receives commands through SCSI target ports or SCSI target/initiator ports.

The service delivery subsystem connects all the SCSI ports in the SCSI domain, providing a mechanism through which application clients and device servers communicate (see 4.6).

The boundaries of a SCSI domain are established by the system implementor, within the constraints of a specific SCSI transport protocol and associated interconnect standards.

4.6 The service delivery subsystem

4.6.1 The service delivery subsystem object

The service delivery subsystem connects SCSI ports (see 3.1.96) and is composed of an interconnect subsystem (see figure 10).




Figure 10 — Service delivery subsystem model

The interconnect subsystem is a set of one or more interconnects that appear to a client or server as a single path for the transfer of requests, responses, and data between SCSI devices.

The service delivery subsystem is assumed to provide error-free transmission of requests and responses between client and server. Although a device driver in a SCSI implementation may perform these transfers through several interactions with its SCSI transport protocol layer, the architecture model portrays each operation, from the viewpoint of the application client, as occurring in one discrete step. The request or response is:

- a) Considered sent when the sender passes it to the SCSI port for transmission;
- b) In transit until delivered; and
- c) Considered received when it has been forwarded to the receiver via the destination SCSI device's SCSI port.

4.6.2 Synchronizing client and server states

One way a client is informed of changes in server state is through the arrival of server responses. Such state changes occur after the server has sent the associated response and possibly before the response has been received by the SCSI initiator device (e.g., the SCSI target device changes state upon processing the **Send Command Complete** procedure call (see 5.4.2), but the SCSI initiator device is not informed of the state change until the **Command Complete Received** SCSI transport protocol service confirmation arrives).

SCSI transport protocols may require the SCSI target device to verify that the response has been received successfully before completing a state change. State changes controlled in this manner are said to be synchronized. Since synchronized state changes are not assumed or required by the architecture model, there may be a time lag between the occurrence of a state change within the SCSI target device and the SCSI initiator device's awareness of that change.

This standard assumes that state synchronization, if required by a SCSI transport protocol standard, is enforced by the service delivery subsystem transparently to the server (i.e., whenever the server invokes a SCSI transport protocol service to return a response as described in 7.8 and 5.4. It is assumed that the SCSI port for such a SCSI transport protocol does not return control to the server until the response has been successfully delivered to the SCSI initiator device).

4.6.3 Request/Response ordering

Request or response transactions are said to be in order if, relative to a given pair of sending and receiving SCSI ports, transactions are delivered in the order they were sent.

A sender may require control over the order in which its requests or responses are presented to the receiver (e.g., the sequence in which requests are received is often important whenever a SCSI initiator device issues a series of commands with the **ORDERED** attribute to a logical unit as described in clause 8). In this case, the order in which these commands are completed, and hence the final state of the logical unit, may depend on the order in which these commands are received. The SCSI initiator device acquires knowledge about the state of pending commands and task management functions and may take action based on the nature and sequence of SCSI target device responses (e.g., if the SCSI initiator device aborts a command whose completion response is in transit and the abort response is received out of order, the SCSI initiator device may incorrectly conclude that no further responses are expected from that command).

The manner in which ordering constraints are established is vendor specific. An implementation may delegate this responsibility to the application client (e.g., the device driver). In-order delivery may be an intrinsic property of the service delivery subsystem or a requirement established by the SCSI transport protocol standard.

The SCSI architecture model assumes in-order delivery to be a property of the service delivery subsystem. This assumption is made to simplify the description of behavior and does not constitute a requirement. This standard makes no assumption about, or places any requirement on the ordering of requests or responses between tasks or task management functions received from different SCSI initiator ports.

4.7 SCSI devices

A SCSI device is a SCSI target device, a SCSI initiator device, or a SCSI target/initiator device.

A SCSI initiator device contains at least one SCSI initiator port and is capable of originating commands and task management requests (see 4.7.1). A SCSI target device contains at least one SCSI target port and is capable of processing commands and task management requests (see 4.7.2). A SCSI target/initiator device contains at least one SCSI target/initiator port and is capable of originating and processing commands and task management requests (see 4.7.3). To be functional, a SCSI domain needs to contain a SCSI target port or a SCSI target/initiator port operating as a SCSI target port and a SCSI initiator port or SCSI target/initiator port operating as a SCSI initiator port.

4.7.1 SCSI initiator device

A SCSI initiator device (see figure 11) contains:

- a) Zero or more initiator device names;
- b) One or more SCSI initiator ports, each containing an:
 - A) Initiator port identifier; and
 - B) Optional initiator port name;
- c) One or more application clients; and
- d) Zero or more application client tasks.

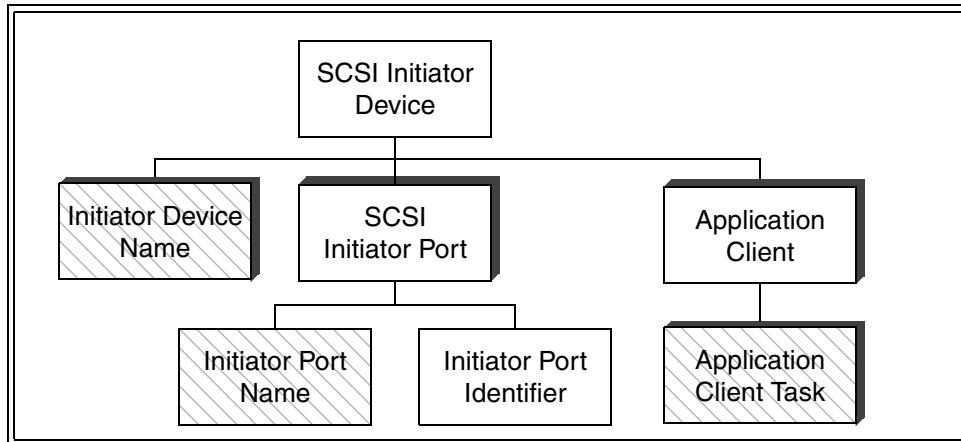


Figure 11 — SCSI initiator device model

The initiator port identifier is a value that is the SCSI port identifier (see 4.7.4) for a SCSI initiator port.

An initiator device name is a name (see 3.1.66) that is a SCSI device name (see 4.7.7) for a SCSI initiator device. For each supported SCSI transport protocol, a SCSI initiator device shall have no more than one (i.e., zero or one) SCSI initiator device name that is not in the SCSI name string format (see SPC-3). A SCSI initiator device shall have no more than one (i.e., zero or one) SCSI initiator device name in the SCSI name string format regardless of the number of SCSI transport protocols supported by the SCSI initiator device. If a SCSI initiator device has a SCSI device name in the SCSI name string format then the SCSI initiator device should have only one SCSI initiator device name. A SCSI transport protocol standard may place additional requirements on initiator device names.

The initiator port name is a name (see 3.1.66) that is the SCSI port name (see 4.7.8) for the SCSI initiator port. A SCSI transport protocol standard may place additional requirements on initiator port names.

Application clients are the sources of commands and task management functions.

An application client task is the source for a single command, series of linked commands, or a single task management function.

4.7.2 SCSI target device

A SCSI target device (see figure 12) contains:

- a) Zero or more target device names;
- b) One or more SCSI target ports, each containing:
 - A) A task router;
 - B) A target port identifier;
 - C) An optional relative port identifier; and
 - D) An optional target port name;
 and
- c) One or more logical units.

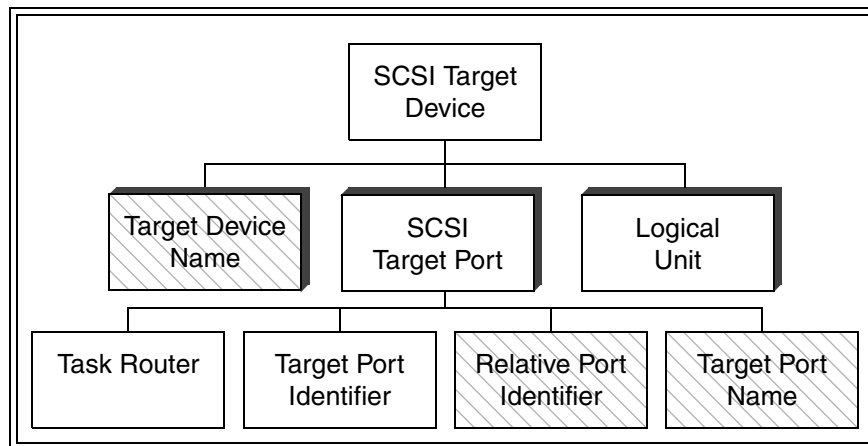


Figure 12 — SCSI target device model

The target port identifier is a value that is a SCSI port identifier (see 4.7.4) for a SCSI target port.

The relative port identifier (see 4.7.5) identifies the SCSI target port relative to other SCSI ports in the SCSI target device.

A target device name is a name (see 3.1.66) that is a SCSI device name (see 4.7.7) for a SCSI target device. For each supported SCSI transport protocol, a SCSI target device shall have no more than one (i.e., zero or one) SCSI target device name that is not in the SCSI name string format (see SPC-3). A SCSI target device shall have no more than one (i.e., zero or one) SCSI target device name in the SCSI name string format regardless of the number of SCSI transport protocols supported by the SCSI target device. If a SCSI target device has a SCSI device name in the SCSI name string format then the SCSI target device should have only one SCSI target device name. A SCSI transport protocol standard may place additional requirements on target device names.

The target port name is a name (see 3.1.66) that is the SCSI port name (see 4.7.8) for the SCSI target port. A SCSI transport protocol standard may place additional requirements on target port names.

The task router routes commands and task management functions between the service delivery subsystem and the appropriate logical unit's task manager (see 4.7.6).

A logical unit is the object to which commands are addressed. One of the logical units within the SCSI target device shall be accessed using the logical unit number zero. See 4.8 for a description of the logical unit.

4.7.3 SCSI target/initiator device

A SCSI target/initiator device contains:

- a) Zero or more target/initiator device names;
- b) One of the following combinations of SCSI ports:
 - A) SCSI target/initiator ports present (see figure 13), with:
 - a) One or more SCSI target/initiator ports each containing:
 - A) A task router;
 - B) A target port identifier;
 - C) An initiator port identifier;
 - D) An optional relative port identifier;
 - E) An optional target port name; and
 - F) An optional initiator port name;
 - b) Zero or more SCSI target ports (see 4.7.2); and
 - c) Zero or more SCSI initiator ports (see 4.7.1);
 - or
 - B) No SCSI target/initiator ports present (see figure 14), with:
 - a) Zero SCSI target/initiator ports;
 - b) One or more SCSI target ports (see 4.7.2); and
 - c) One or more SCSI initiator ports (see 4.7.1);
- c) One or more logical units; and
- d) One or more application clients (see 4.7.1).

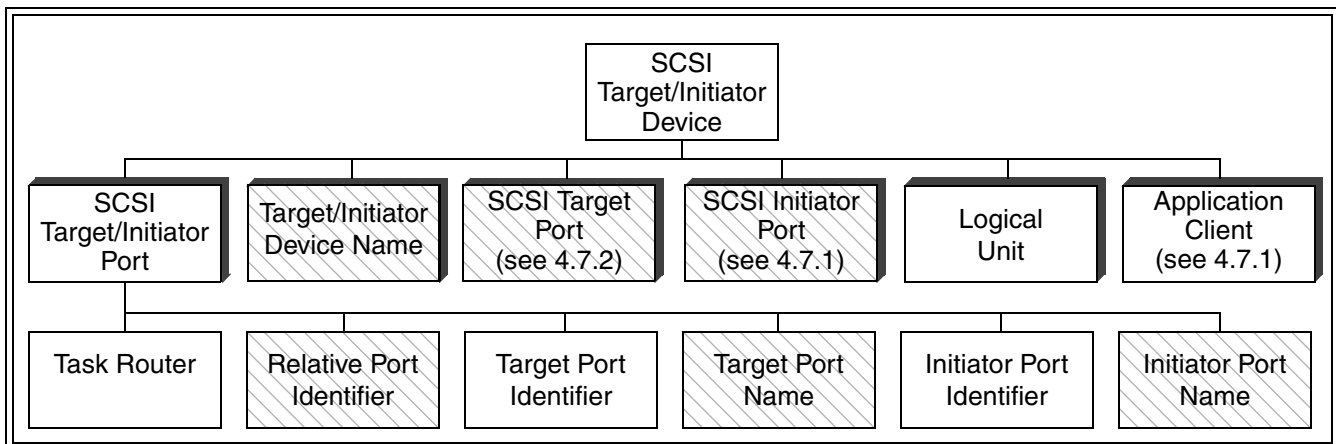


Figure 13 — SCSI target/initiator device with SCSI target/initiator ports model

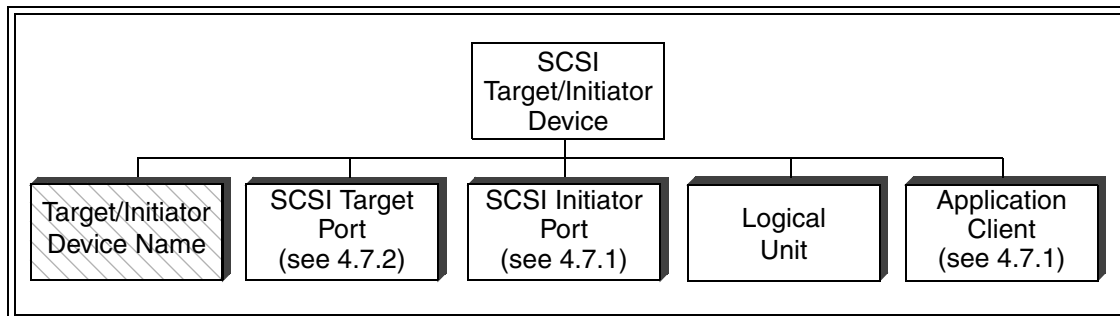


Figure 14 — SCSI target/initiator device without SCSI target/initiator ports model

The target port identifier and the initiator port identifier are values containing a SCSI port identifier (see 4.7.4) for a SCSI target/initiator port. The target port identifier and the initiator port identifier may or may not be identical.

The relative port identifier (see 4.7.5) identifies the SCSI target port relative to other SCSI ports in the SCSI target device.

A target/initiator device name is a name (see 3.1.66) that is a SCSI device name (see 4.7.7) for a SCSI target/initiator device. For each supported SCSI transport protocol, a SCSI target/initiator device shall have no more than one (i.e., zero or one) SCSI target/initiator device name that is not in the SCSI name string format (see SPC-3). A SCSI target/initiator device shall have no more than one (i.e., zero or one) SCSI target/initiator device name in the SCSI name string format regardless of the number of SCSI transport protocols supported by the SCSI target/initiator device. If a SCSI target/initiator device has a SCSI device name in the SCSI name string format then the SCSI target/initiator device should have only one SCSI target/initiator device name. A SCSI transport protocol standard may place additional requirements on target/initiator device names.

The target port name and the initiator port name are names (see 3.1.66) that are the SCSI port name (see 4.7.8) for the target/initiator port when operating as a SCSI target port and SCSI initiator port, respectively. The target port name and the initiator port name may or may not be identical. A SCSI transport protocol standard may place additional requirements on target port names and initiator port names.

When the SCSI target/initiator device is operating as a SCSI target device the task router routes the commands and task management functions between the service delivery subsystem and the appropriate logical unit (see 4.7.6).

A logical unit is the object to which commands are sent. One of the logical units within the SCSI target/initiator device shall be accessed using the logical unit number zero. See 4.8 for a description of the logical unit.

When the SCSI target/initiator device is operating as a SCSI initiator device an application client is the source of commands and task management functions.

4.7.4 SCSI port identifier

The SCSI port identifier is equivalent to SCSI identifier. The SCSI port identifier object represents either an initiator port identifier for a SCSI initiator port, or a target port identifier for a SCSI target port. SCSI port identifier is used when either a SCSI initiator port or SCSI target port is applicable or when other context in the description identifies the SCSI initiator port or SCSI target port usage.

4.7.5 Relative port identifier

A SCSI target device or a SCSI target/initiator device may assign each of its SCSI ports a unique relative port identifier from 1 to 65 535. SCSI target ports, SCSI target/initiator ports, and SCSI initiator ports share the same number space.

Relative port identifiers may be retrieved through the Device Identification VPD page (see SPC-3) and the SCSI Ports VPD page (see SPC-3).

The relative port identifiers are not required to be contiguous. The relative port identifier for a SCSI port shall not change once assigned unless physical reconfiguration of the SCSI target device occurs.

4.7.6 SCSI task router

The task router routes tasks and task management functions to the selected logical unit. Any task that is sent to a logical unit that is not known to the task router is handled as described in 5.9.4. Any task management function that is not sent to a specific logical unit shall be broadcast to all logical units known to the task router.

4.7.7 SCSI device name

A SCSI device name is an optional name (see 3.1.66) for a SCSI device that is world wide unique within the SCSI transport protocol of each SCSI domain in which the SCSI device has SCSI ports. For each supported SCSI transport protocol, a SCSI device shall have no more than one (i.e., zero or one) SCSI device name that is not in the SCSI name string format (see SPC-3). A SCSI device shall have no more than one (i.e., zero or one) SCSI device name in the SCSI name string format regardless of the number of SCSI transport protocols supported by the SCSI device. If a SCSI device has a SCSI device name in the SCSI name string format then the SCSI device should have only one SCSI device name. A SCSI device name shall never change and may be used to persistently identify a SCSI device in contexts where specific references to port names or port identifiers is not required.

A SCSI transport protocol standard may require that a SCSI device include a SCSI device name if the SCSI device has SCSI ports in a SCSI domain of that SCSI transport protocol. The SCSI device name may be made available to other SCSI devices or SCSI ports in a given SCSI domain in SCSI transport protocol specific ways.

4.7.8 SCSI port name

A SCSI port name is an optional name (see 3.1.66) of a SCSI port that is world wide unique within the SCSI transport protocol of the SCSI domain of that SCSI port. A SCSI port may have at most one name. A SCSI port name shall never change and may be used to persistently identify a SCSI initiator port or SCSI target port in contexts similar to those where a SCSI port identifier (see 4.7.4) may be used.

A SCSI transport protocol standard may require that a SCSI port include a SCSI port name if the SCSI port is in a SCSI domain of that SCSI transport protocol. The SCSI port name may be made available to other SCSI devices or SCSI ports in the given SCSI domain in SCSI transport protocol specific ways.

4.8 Logical units

A logical unit (see figure 15) contains:

- a) A logical unit number;
 - A) If access controls (see SPC-3) are not in effect, one logical unit number per logical unit; or
 - B) If access controls are in effect, one logical unit number per SCSI initiator port that has access rights plus one default logical unit number per logical unit;
- b) One or more logical unit names;
- c) A device server;
- d) A task manager; and
- e) One or more task sets each of which may contain zero or more tasks.

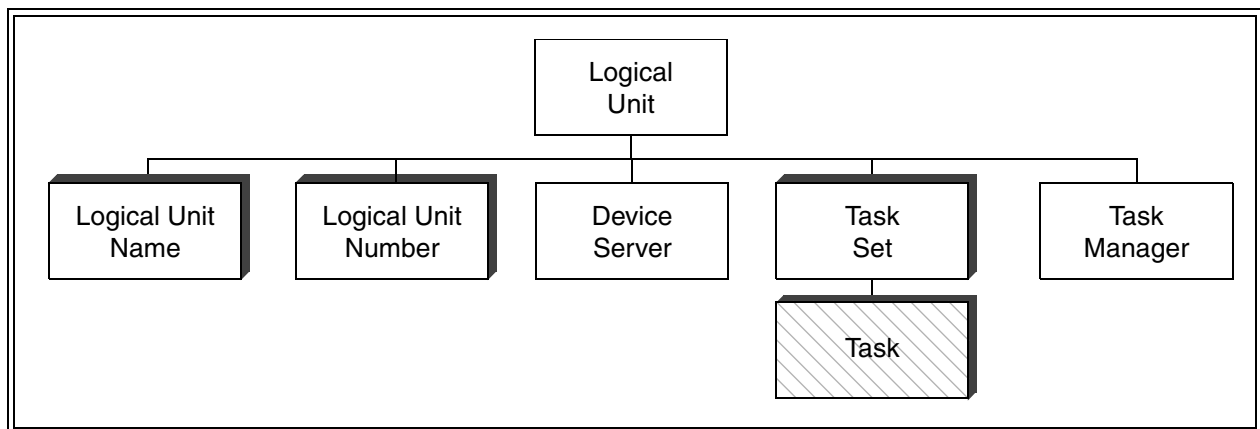


Figure 15 — Logical unit model

A logical unit number is a field (see 4.9) containing 64 bits that identifies the logical unit within a SCSI target device when accessed by a SCSI target port. If any logical unit within the scope of a SCSI target device includes dependent logical units in its composition, then all logical unit numbers within the scope of the SCSI target device shall have the format described in 4.9.4. Otherwise, the logical unit numbers should have the format described in 4.9.3.

A logical unit name is a name (see 3.1.66) for a logical unit that is not a well known logical unit. A logical unit name shall be world wide unique. A logical unit name shall never change and may be used to persistently identify a logical unit.

A SCSI transport protocol standard may require that a logical unit include a logical unit name if the logical unit is accessed through a SCSI domain defined by that SCSI transport protocol.

The device server is the object that processes the operations requested by the received commands.

The task manager controls the sequencing of one or more tasks within a logical unit. The task manager also carries out the task management functions specified in clause 7. There is one task manager per logical unit.

The order in which task management requests are processed is not specified by this standard. This standard does not require in-order delivery of such requests, as defined in 4.6.3, or processing by the task manager in the order received. To guarantee the processing order of task management requests referencing a specific logical unit, an application client should not have more than one such request pending to that logical unit.

A task set is composed of zero or more tasks (see 4.11).

The interactions among the tasks in a task set are determined by the requirements for task set management specified in clause 8 and the ACA requirements specified in 5.9.1. The number of task sets per logical unit and the boundaries between task sets are governed by the TST field in the Control mode page (see SPC-3).

4.9 Logical unit numbers

4.9.1 Logical unit numbers overview

All logical unit number formats described in this standard are hierarchical in structure even when only a single level in that hierarchy is used. The HISUP bit shall be set to one in the standard INQUIRY data (see SPC-3) when any logical unit number format described in this standard is used. Non-hierarchical formats are outside the scope of this standard.

4.9.2 LUN 0 address

All SCSI devices shall accept LUN 0 as a valid address. For SCSI devices that support the hierarchical addressing model the LUN 0 shall be the logical unit that an application client addresses to determine information about the SCSI target device and the logical units contained within the SCSI target device.

To address the LUN 0 of a SCSI device the peripheral device address method shall be used.

4.9.3 Single level logical unit number structure

When the single level subset format is used, the HISUP bit shall be set to one in the standard INQUIRY data (see SPC-3) returned by the logical unit with the logical unit number zero.

Table 1 describes a single level subset of the format described in 4.14 for logical unit numbers 255 and below.

Table 1 — Single level logical unit number structure for logical unit numbers 255 and below

Bit Byte	7	6	5	4	3	2	1	0
0	ADDRESS METHOD (00b)		BUS IDENTIFIER (00h)					
1	SINGLE LEVEL LUN (00h to FFh, inclusive)							
2	(MSB)	Null second level LUN (0000h)						(LSB)
3								
4	(MSB)	Null third level LUN (0000h)						(LSB)
5								
6	(MSB)	Null fourth level LUN (0000h)						(LSB)
7								

All logical unit number structure fields shall be zero except the SINGLE LEVEL LUN field (see table 1). The value in the SINGLE LEVEL LUN field shall be between 0 and 255, inclusive. The 00b in the ADDRESS METHOD field specifies peripheral device addressing (see 4.9.4) and the 00h in the BUS IDENTIFIER field specifies the current level (see 4.9.6).

Table 2 describes a single level subset of the format described in 4.14 for logical unit numbers 16 383 and below.

Table 2 — Single level logical unit number structure for logical unit numbers 16 383 and below

Bit Byte	7	6	5	4	3	2	1	0
0	ADDRESS METHOD (01b)		(MSB)					
1	SINGLE LEVEL LUN (0000h to 3FFFh, inclusive)						(LSB)	
2	(MSB)		Null second level LUN (0000h)					
3							(LSB)	
4	(MSB)		Null third level LUN (0000h)					
5							(LSB)	
6	(MSB)		Null fourth level LUN (0000h)					
7							(LSB)	

All logical unit number structure fields shall be zero except the SINGLE LEVEL LUN field (see table 2). The value in the SINGLE LEVEL LUN field shall be between 0 and 16 383, inclusive. The 01b in the ADDRESS METHOD field specifies flat space addressing (see 4.9.7) at the current level.

If a SCSI target device contains 256 or fewer logical units, none of which are dependent logical units (see 4.14) or extended addressing logical units (see 4.9.8), then its logical units should be numbered 255 and below.

If a SCSI target device contains 16 384 or fewer logical units, none of which are dependent logical units or extended addressing logical units, then its logical units should be numbered 16 383 and below.

- █ Except for dependent logical units and extended addressing logical units, logical unit numbers that are greater than 255 shall have the format shown in table 2. Except for dependent logical units and extended addressing logical units, logical unit numbers that are less than 256 should have the format shown in table 1 but may have the format shown in table 2.

4.9.4 Eight byte logical unit number structure

The eight byte logical unit number structure (see table 4) allows up to four levels of SCSI devices to be addressed under a single SCSI target device. Each level shall use byte 0 and byte 1 to define the address and/or location of the SCSI device to be addressed on that level.

If the logical unit number specifies that the command is to be relayed to the next layer then the current layer shall use byte 0 and byte 1 of the eight byte logical unit number structure to determine the address of the SCSI device to which the command is to be sent. When the command is sent to the SCSI target device the eight byte logical unit number structure that was received shall be adjusted to create a new eight byte logical unit number structure (see table 3 and figure 16).

SCSI devices shall keep track of the addressing information necessary to transmit information back through all intervening layers to the task's originating SCSI initiator port.




Figure 16 — Eight Byte logical unit number structure adjustments

Table 3 — Eight byte logical unit number structure adjustments

Byte position		
Old		New
0 & 1	Moves to	Not Used
2 & 3	Moves to	0 & 1
4 & 5	Moves to	2 & 3
6 & 7	Moves to	4 & 5
N/A	zero fill	6 & 7

The eight byte logical unit number structure requirements as viewed from the application client are shown in table 4.

Table 4 — Eight Byte logical unit number structure

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)	FIRST LEVEL ADDRESSING						(LSB)
1		SECOND LEVEL ADDRESSING						(LSB)
2	(MSB)	THIRD LEVEL ADDRESSING						(LSB)
3		FOURTH LEVEL ADDRESSING						(LSB)
4	(MSB)							(LSB)
5								(LSB)
6	(MSB)							(LSB)
7								(LSB)

The FIRST LEVEL ADDRESSING field specifies the first level address of a SCSI device. See table 5 for a definition of the FIRST LEVEL ADDRESSING field.

The SECOND LEVEL ADDRESSING field specifies the second level address of a SCSI device. See table 5 for a definition of the SECOND LEVEL ADDRESSING field.

The THIRD LEVEL ADDRESSING field specifies the third level address of a SCSI device. See table 5 for a definition of the THIRD LEVEL ADDRESSING field.

The FOURTH LEVEL ADDRESSING field specifies the fourth level address of a SCSI device. See table 5 for a definition of the FOURTH LEVEL ADDRESSING field.

The SCSI device pointed to in the FIRST LEVEL ADDRESSING, SECOND LEVEL ADDRESSING, THIRD LEVEL ADDRESSING, and FOURTH LEVEL ADDRESSING fields may be any physical or logical device addressable by an application client.

Table 5 — Format of addressing fields

Bit Byte	7	6	5	4	3	2	1	0	
n-1	ADDRESS METHOD		(MSB)						(LSB)
n	ADDRESS METHOD SPECIFIC							(LSB)	

The ADDRESS METHOD field defines the contents of the ADDRESS METHOD SPECIFIC field. See table 6 for the address methods defined for the ADDRESS METHOD field. The ADDRESS METHOD field only defines address methods for entities that are directly addressable by an application client.

Table 6 — ADDRESS METHOD field values

Code	Description	Reference
10b	Logical unit addressing method	4.9.5
00b	Peripheral device addressing method	4.9.6
01b	Flat space addressing method	4.9.7
11b	Extended logical unit addressing method	4.9.8

4.9.5 Logical unit addressing method

If the logical unit addressing method is selected the SCSI device should relay the received command to the addressed dependent logical unit. Any command that is not relayed to a dependent logical unit shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

NOTE 1 - A SCSI device may filter (i.e., not relay) commands to prevent commands with deleterious effects from reaching a dependent logical unit (e.g., a WRITE command directed to a logical unit that is participating in a RAID volume).

See table 7 for the definition of the ADDRESS METHOD SPECIFIC field used when the logical unit addressing method is selected.

Table 7 — Logical unit addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	1	0	TARGET					
n	BUS NUMBER			LUN				

The TARGET field, BUS NUMBER field, and LUN field address the logical unit to which the received command shall be relayed. The command shall be relayed to the logical unit specified by the LUN field within the SCSI target device specified by the TARGET field located on the bus specified by the BUS NUMBER field. The SCSI target device information in the TARGET field may be a target port identifier (see 4.7.2) or it may be a mapped representation of a target port identifier, when the range of possible target port identifiers is too large to fit in the TARGET field.

NOTE 2 - The value of target port identifiers within the TARGET field are defined by individual standards. (e.g., SCSI Parallel Interface -2 standard defines target port identifiers to be in the range 0 to 7, 0 to 15, and 0 to 31).

4.9.6 Peripheral device addressing method

If the peripheral device addressing method is selected, the SCSI device should relay the received command to the addressed dependent logical unit. Any command that is not relayed to a dependent logical unit shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

NOTE 3 - A SCSI device may filter (i.e., not relay) commands to prevent commands with deleterious effects from reaching a dependent logical unit (e.g., a WRITE command directed to a logical unit that is participating in a RAID volume).

See table 8 for the definition of the ADDRESS METHOD SPECIFIC field used when the peripheral device addressing method is selected.

Table 8 — Peripheral device addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	0	0	BUS IDENTIFIER					
n	TARGET/LUN							

The BUS IDENTIFIER field identifies the bus or path that the SCSI device shall use to relay the received command. The BUS IDENTIFIER field may use the same value encoding as the BUS NUMBER field (see 4.9.5). However, bus identifier zero shall specify that the command is to be relayed to a logical unit within the SCSI device at the current level.

The TARGET/LUN field specifies the address of the peripheral device to which the SCSI device shall relay the received command. The meaning and usage of the TARGET/LUN field depends on whether the BUS IDENTIFIER field contains zero.

A BUS IDENTIFIER field of zero specifies a logical unit at the current level. This representation of a logical unit may be used either when the SCSI device at the current level does not use hierarchical addressing for assigning LUNs to entities or when the SCSI device at the current level includes entities that need LUNs but are not attached to SCSI buses (e.g., fans, cache, and controllers). When the BUS IDENTIFIER field contains zero, the command shall be relayed to the current level logical unit specified by the TARGET/LUN field within or joined to the current level SCSI device.

A BUS IDENTIFIER field greater than zero represents a SCSI domain that connects a group of SCSI devices to the current level SCSI device. Each SCSI domain shall be assigned a unique bus identifier number from 1 to 63. These bus identifiers shall be used in the BUS IDENTIFIER field when assigning addresses to peripheral devices attached to the SCSI domains. When the BUS IDENTIFIER field is greater than zero, the command shall be relayed to the logical unit with the logical unit number zero within the SCSI target device specified in the TARGET/LUN field located in the SCSI domain specified by the BUS IDENTIFIER field. The SCSI target device information in the TARGET/LUN field may be a target port identifier (see 4.7.2) or it may be a mapped representation of a target port identifier, when the range of possible target port identifiers is too large to fit in the TARGET/LUN field.

The SCSI device located within the current level shall be addressed by a BUS IDENTIFIER field and a TARGET/LUN field of all zeros, also known as LUN 0 (see 4.9.2).

4.9.7 Flat space addressing method

All commands are allowed when the flat space addressing method is used, however, the addressed logical unit is not required to support all commands. Any command that is not supported shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID COMMAND OPERATION CODE.

In the response to an INQUIRY command, the addressed logical unit shall return a valid SCSI peripheral device type (e.g., direct access device, streaming device).

See table 9 for the definition of the ADDRESS METHOD SPECIFIC field used when the flat space addressing method is selected.

Table 9 — Flat space addressing

Bit Byte	7	6	5	4	3	2	1	0
n-1	0	1	(MSB)					
n	LUN						(LSB)	

The LUN field specifies the address of the logical unit to which the current level shall direct the received command.

4.9.8 Extended logical unit addressing

Extended logical unit addressing builds on the formats defined for dependent logical units (see 4.14) but may be used by SCSI devices having single level logical unit structure. In dependent logical unit addressing, the logical unit information at each level fits in exactly two bytes. Extended logical unit addresses have sizes of two bytes, four bytes, six bytes, or eight bytes.

Extended logical units are identified by the ADDRESS METHOD field (see table 6 in 4.14) in the same manner as is the case for dependent logical units. An ADDRESS METHOD field value of 11b specifies the extended logical unit addressing method.

See table 10 for the definition of the ADDRESS METHOD SPECIFIC field used when the extended logical unit addressing method is selected.

Table 10 — Extended logical unit addressing

Bit Byte	7	6	5	4	3	2	1	0
n	1	1	LENGTH		EXTENDED ADDRESS METHOD			
m	EXTENDED ADDRESS METHOD SPECIFIC							

The LENGTH field (see table 11) specifies the length of the EXTENDED ADDRESS METHOD SPECIFIC field.

Table 11 — LENGTH field values

Value	Length of the EXTENDED ADDRESS METHOD SPECIFIC Field	Reference
00b	One byte	table 12
01b	Three bytes	table 13
10b	Five bytes	table 14
11b	Seven bytes	table 15

Table 12, table 13, table 14, and table 15 show the four extended logical unit addressing formats.

Table 12 — Two byte extended logical unit addressing format

Bit Byte	7	6	5	4	3	2	1	0
n	1	1	LENGTH (00b)		EXTENDED ADDRESS METHOD			
n+1	EXTENDED ADDRESS METHOD SPECIFIC							

Table 13 — Four byte extended logical unit addressing format

Bit Byte	7	6	5	4	3	2	1	0
n	1	1	LENGTH (01b)		EXTENDED ADDRESS METHOD			
n+1	(MSB) _____ EXTENDED ADDRESS METHOD SPECIFIC _____ (LSB)							
n+3								

Table 14 — Six byte extended logical unit addressing format

Bit Byte	7	6	5	4	3	2	1	0
n	1	1	LENGTH (10b)		EXTENDED ADDRESS METHOD			
n+1	(MSB) _____ EXTENDED ADDRESS METHOD SPECIFIC _____ (LSB)							
n+5								

Table 15 — Eight byte extended logical unit addressing format

Bit Byte	7	6	5	4	3	2	1	0
0	1	1	LENGTH (11b)		EXTENDED ADDRESS METHOD			
1	(MSB) _____ EXTENDED ADDRESS METHOD SPECIFIC _____ (LSB)							
7								

The EXTENDED ADDRESS METHOD field combined with the LENGTH field (see table 16) specifies the type and size of extended logical unit address found in the EXTENDED ADDRESS METHOD SPECIFIC field.

Table 16 — Logical unit extended address methods

EXTENDED ADDRESS METHOD Code	LENGTH Code(s)	Description	Reference
0h	00b - 11b	Reserved	
1h	00b	Well known logical unit	4.9.9
1h	01b - 11b	Reserved	
2h - Eh	00b - 11b	Reserved	
Fh	00b - 10b	Reserved	
Fh	11b	Logical unit not specified	4.9.10

4.9.9 Well known logical unit addressing

A SCSI target device may support zero or more well known logical units. A single SCSI target device shall only support one instance of each supported well known logical unit. All well known logical units within a SCSI target device shall be accessible from all SCSI target ports contained within the SCSI target device.

See table 17 for the definition of the EXTENDED ADDRESS METHOD SPECIFIC field used when the well known logical unit extended address method is selected.

Table 17 — Well known logical unit extended address format

Bit Byte	7	6	5	4	3	2	1	0
n	1	1	LENGTH (00b)		Well known logical unit (1h)			
n+1	W-LUN							

The W-LUN field specifies well known logical unit to be addressed (see SPC-3).

4.9.10 Logical unit not specified addressing

A logical unit not specified extended address method LUN shall be used to indicate that no logical unit of any kind is specified.

See table 18 for the format of a LUN using the logical unit not specified extended address method.

Table 18 — Logical unit not specified extended address method format

Bit Byte	7	6	5	4	3	2	1	0
0	1	1	LENGTH (11b)		Logical unit not specified (Fh)			
1	FFh							
2	FFh							
3	FFh							
4	FFh							
5	FFh							
6	FFh							
7	FFh							

4.10 Well known logical units

A well known logical unit is identical to a logical unit (see 4.8) except as defined in this subclause.

Well known logical units are addressed using the well known logical unit addressing method of extended logical unit addressing (see 4.9.8). Each well known logical unit has a well known logical unit number (W-LUN). W-LUN values are defined in SPC-3.

If a SCSI target device receives a W-LUN and the well known logical unit specified by the W-LUN does not exist, a task manager shall follow the rules for selection of invalid logical units described in 5.9.4.

If a well known logical unit is supported within a SCSI target device then that logical unit shall support all the commands defined for it.

Access to well known logical units shall not be affected by access controls.

All well known logical units:

- a) Shall not have logical unit names; and
- b) Shall identify themselves using the SCSI target device names of the SCSI target device in which they are contained.

A SCSI target device may have more than one SCSI target device name if the SCSI target device supports multiple SCSI transport protocols.

The name of the well known logical unit may be determined by issuing an INQUIRY command requesting the Device Identification VPD page (see SPC-3).

4.11 Tasks and task tags

A task is represented by an I_T_L_Q nexus (see 4.12) and is composed of:

- a) A definition of the work to be performed by the logical unit in the form of a command or a group of linked commands; and
- b) A task attribute (see 8.6) that allows the application client to specify processing relationships between various tasks in the task set.

The I_T_L_Q nexus representing a task includes a task tag, allowing many uniquely identified tagged tasks to be present in a single task set. A task tag is a value that is composed of up to 64 bits.

A SCSI initiator device assigns task tag values for each I_T_L_Q nexus in a way that ensures that the nexus uniqueness requirements stated in this subclause are met. Transport protocols may define additional restrictions on task tag assignment (e.g., restricting task tag length, requiring task tags to be unique per I_T nexus or per I_T_L nexus, or sharing task tag values with other uses such as task management functions).

An I_T_L_Q nexus that is in use (i.e., during the interval bounded by the events specified in 5.5) shall be unique as seen by the SCSI initiator port originating the command and the logical unit to which the command was addressed, otherwise an overlapped command condition exists (see 5.9.3). An I_T_L_Q nexus is unique if one or more of its components is unique within the specified time interval.

A SCSI initiator device shall not create more than one task from a specific SCSI initiator port having identical values for the target port identifier, logical unit number, and task tag.

4.12 The nexus object

The nexus object represents a relationship between a SCSI initiator port, a SCSI target port, optionally a logical unit, and optionally a task.

The nexus object may refer to any one or all of the following relationships:

- a) One SCSI initiator port to one SCSI target port (an I_T nexus);
- b) One SCSI initiator port to one SCSI target port to one logical unit (an I_T_L nexus);
- c) One SCSI initiator port to one SCSI target port to one logical unit to one task (an I_T_L_Q nexus); or
- d) Either an I_T_L nexus or an I_T_L_Q nexus (denoted as an I_T_L_x nexus).

Table 19 maps the nexus object to other identifier objects.

Table 19 — Mapping nexus to SAM-2 identifiers

Nexus	Identifiers contained in nexus	Reference
I_T	Initiator Port Identifier	4.7.1
	Target Port Identifier	4.7.2
I_T_L	Initiator Port Identifier	4.7.1
	Target Port Identifier	4.7.2
	Logical Unit Number	4.8
I_T_L_Q	Initiator Port Identifier	4.7.1
	Target Port Identifier	4.7.2
	Logical Unit Number	4.8
	Task Tag	4.11

4.13 SCSI ports

4.13.1 SCSI port configurations

A SCSI device may contain only SCSI target ports, only SCSI initiator ports, only SCSI target/initiator ports or any combination of ports. Some of the port configurations possible for a SCSI device are shown in figure 17.

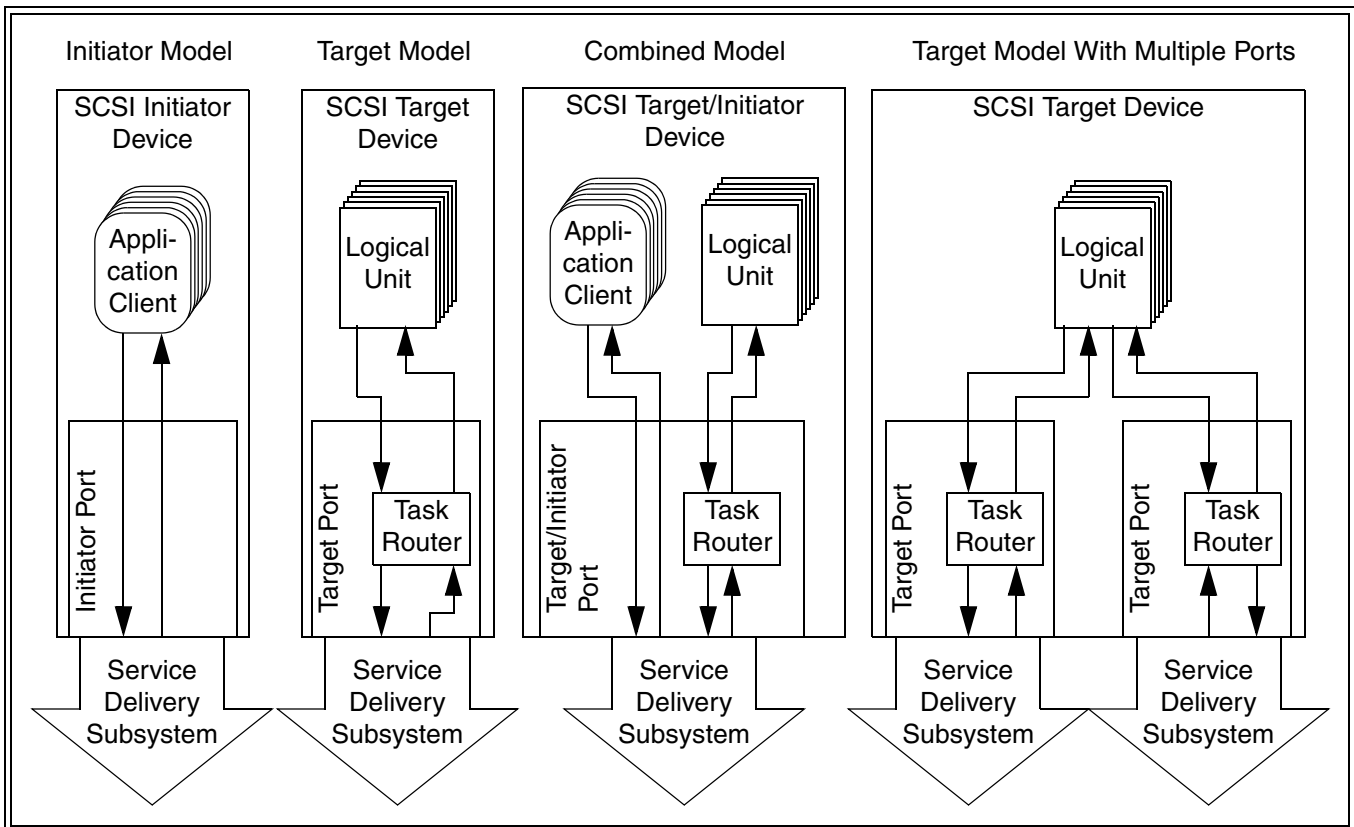


Figure 17 — SCSI device functional models

A SCSI target/initiator device is referred to by the role its port takes when it participates in an I/O operation. When a SCSI target/initiator device receives commands or task management functions, the SCSI target/initiator device takes on the characteristics of and is referred to as a SCSI target device. When a SCSI target/initiator device issues commands or task management functions, the SCSI target/initiator device takes on the characteristics of and is referred to as a SCSI initiator device.

4.13.2 SCSI devices with multiple ports

The model for a SCSI device with multiple ports is a single SCSI target device (see 4.7.2), SCSI initiator device (see 4.7.1), or SCSI target/initiator device (see 4.7.3) with multiple ports. Similarly, a single SCSI target port or SCSI initiator port may respond to multiple SCSI identifiers, with the model for such a SCSI port being one of multiple SCSI target ports or SCSI initiator ports (i.e., one for each SCSI identifier).

The SCSI identifiers representing the ports shall meet the requirements for initiator port identifiers (see 4.7.1) or target port identifiers (see 4.7.2) or both. SCSI target/initiator devices with multiple ports implement both target and initiator models and combine the SCSI target/initiator port structures in vendor specific ways that meet product requirements while maintaining the model for SCSI devices with multiple ports for the target and initiator functions

performed by the product. How a multiple port SCSI device is viewed by counterpart SCSI devices in the SCSI domain also depends on whether a SCSI initiator port is examining a SCSI target port or SCSI target/initiator port, or a SCSI target port is servicing a SCSI initiator port or SCSI target/initiator port. The structures and views of SCSI devices are asymmetric for SCSI target ports and SCSI initiator ports.

4.13.3 Multiple port target SCSI device structure

Figure 18 shows the structure of a SCSI target device with multiple SCSI target ports. Each SCSI target port consists of a task router that is shared by a collection of logical units. Each logical unit contains a single task manager and a device server.

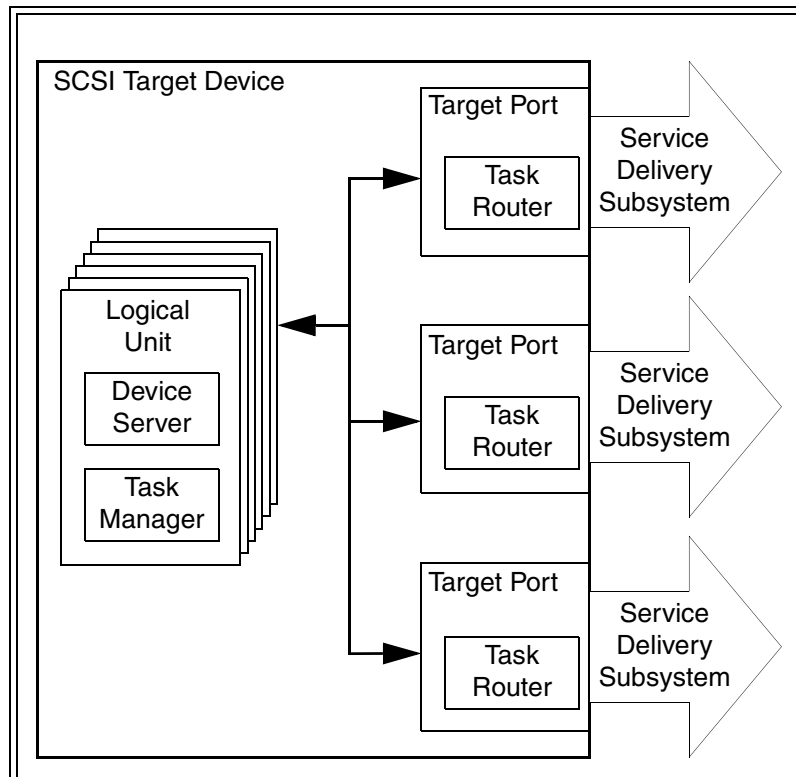


Figure 18 — Multiple port target SCSI device structure model

Two-way communications shall be possible between all logical units and all SCSI target ports, however, communications between any logical unit and any SCSI target port may be inactive. Two-way communications shall be available between each task manager and all task routers. Each SCSI target port shall accept commands sent to LUN 0 and the task router shall route them to a device server for processing. The REPORT LUNS commands (see SPC-3) shall be accepted by the logical unit with the logical unit number zero from any SCSI target port and shall return the logical unit inventory available via that SCSI target port. The availability of the same logical unit through multiple SCSI target ports is discovered by matching SCSI port name or identifier values in the INQUIRY command Device Identification VPD page (see SPC-3).

4.13.4 Multiple port initiator SCSI device structure

Figure 19 shows the structure of a SCSI initiator device with multiple SCSI initiator ports. Each SCSI initiator port is shared by a collection of application clients.

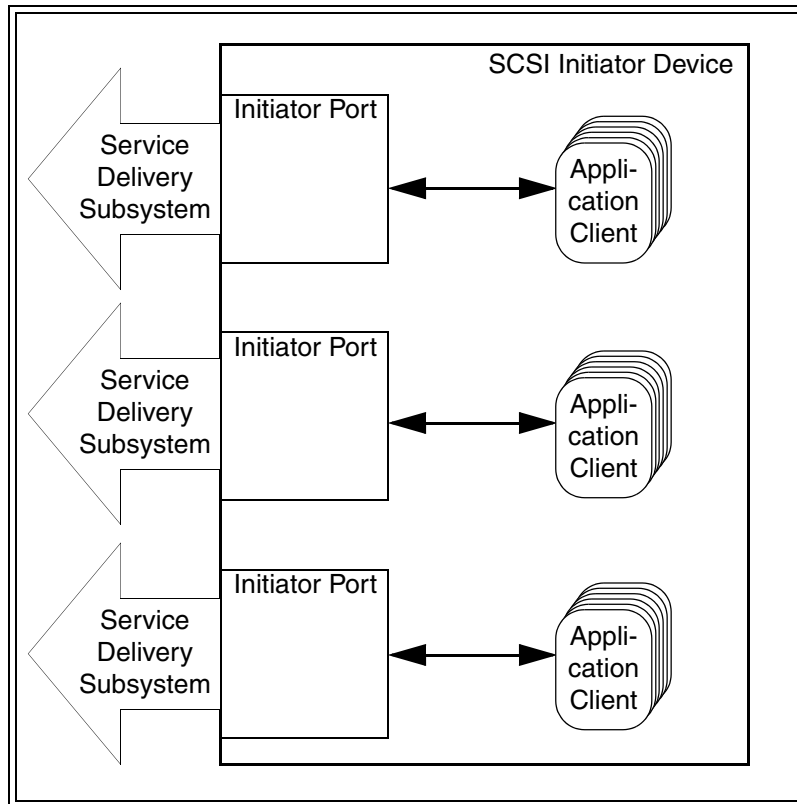


Figure 19 — Multiple port SCSI initiator device structure model

Two-way communications shall be possible between an application client and its associated SCSI initiator port. This standard does not specify or require the definition of any mechanisms by which a SCSI target device would have the ability to discover that it is communicating with multiple ports on a single SCSI initiator device. In those SCSI transport protocols where such mechanisms are defined, they shall not have any effect on how commands are processed (e.g., reservations shall be handled as if no such mechanisms exist).

4.13.5 Multiple port target/initiator SCSI device structure

Figure 20 shows the structure of a SCSI target/initiator device with multiple SCSI target/initiator ports. Each SCSI target/initiator port consists of a task router and is shared by a collection of logical units and application clients. Each logical unit contains a single task manager and a device server.

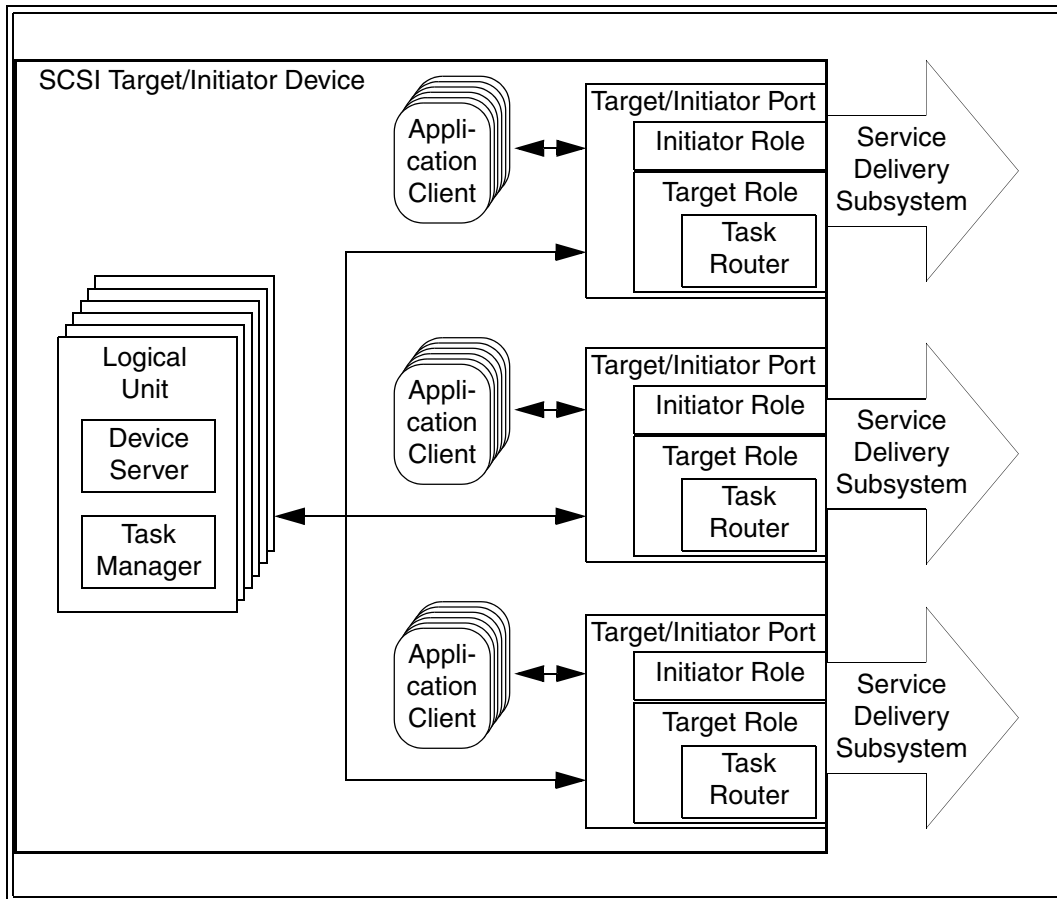


Figure 20 — Multiple port target/initiator SCSI device structure model

Two-way communications shall be possible between all logical units and all SCSI target/initiator ports, however, communications between any logical unit and any SCSI target/initiator port may be inactive. Two-way communications shall be possible between an application client and its associated SCSI target/initiator port. Each SCSI target/initiator port shall accept commands sent to LUN 0 and the task router shall route them to a device server for processing. The REPORT LUNS commands (see SPC-3) shall be accepted by the logical unit with the logical unit number zero from any SCSI target/initiator port and shall return the logical unit inventory available via that SCSI target/initiator port. The availability of the same logical unit through multiple SCSI target/initiator ports is discovered by matching SCSI port name or identifier values in the INQUIRY command Device Identification VPD page (see SPC-3).

This standard does not specify or require the definition of any mechanisms by which a SCSI target device would have the ability to discover that it is communicating with multiple SCSI initiator ports on a single SCSI target/initiator device. In those SCSI transport protocols where such mechanisms are defined, they shall not have any effect on how commands are processed (e.g., reservations shall be handled as if no such mechanisms exist).

4.13.6 SCSI initiator device view of a multiple port SCSI target device

A SCSI target device may be connected to multiple domains such that a SCSI initiator port is only able to communicate with its logical units using a single SCSI target port. However, SCSI target devices with multiple SCSI ports may be configured where application clients have the ability to discover that one or more logical units are accessible via multiple SCSI target ports. Figure 21 and figure 22 show two examples of such configurations.

Figure 21 shows a SCSI target device with multiple SCSI target ports participating in a single SCSI domain with two SCSI initiator devices. There are three SCSI devices, one of which has two SCSI target ports, and two of which have one SCSI initiator port each. There are two target port identifiers and two initiator port identifiers in this SCSI domain. Using the INQUIRY command Device Identification VPD page (see SPC-3), the application clients in each of the SCSI initiator devices have the ability to discover the logical units in the SCSI target devices are accessible via multiple target port identifiers (i.e., SCSI target ports) and map the configuration of the SCSI target devices.

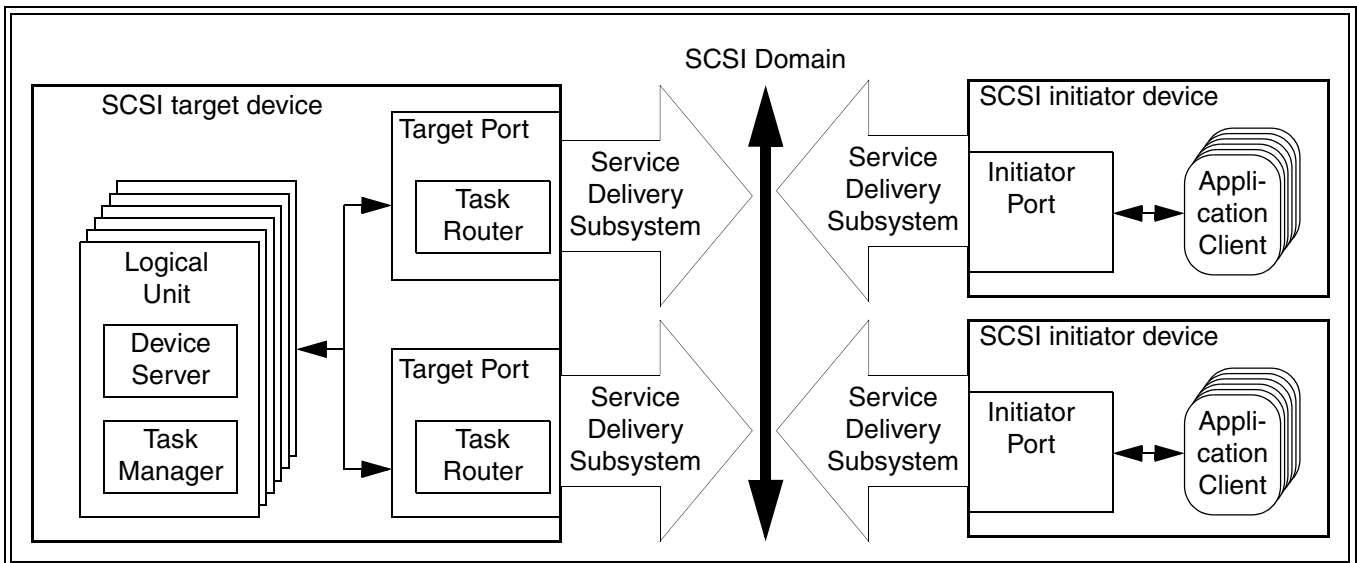


Figure 21 — SCSI target device configured in a single SCSI domain

Figure 22 shows a SCSI target device with multiple SCSI target ports participating in two SCSI domains and a SCSI initiator device with multiple SCSI initiator ports participating in the same two SCSI domains. There is one SCSI target device with two SCSI target ports and one SCSI initiator device with two SCSI initiator ports. There is one target port identifier and one initiator port identifier in each of the two SCSI domains. Using the INQUIRY command Device Identification VPD page (see SPC-3), the application clients in the SCSI initiator device have the ability to discover that logical units in the SCSI target device are accessible via multiple ports and map the configuration. However, application clients may not be able to distinguish between the configuration shown in figure 22 and the configuration shown in figure 23.

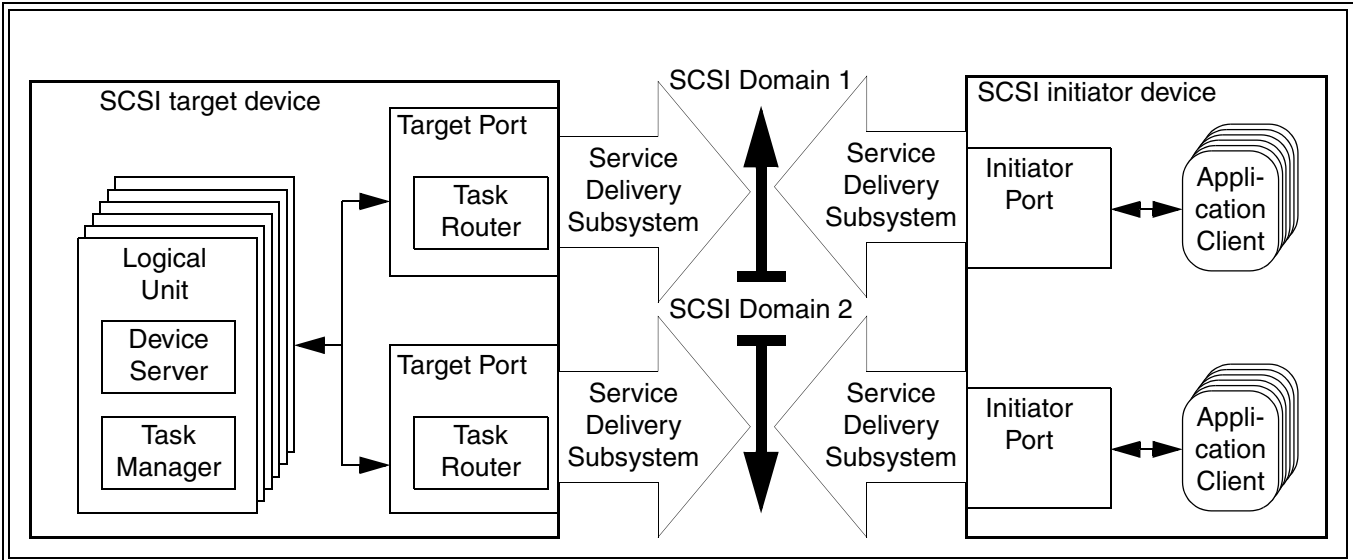


Figure 22 — SCSI target device configured in multiple SCSI domains

Figure 23 shows the same configuration as figure 22 except that the two SCSI domains have been replaced by a single SCSI domain.

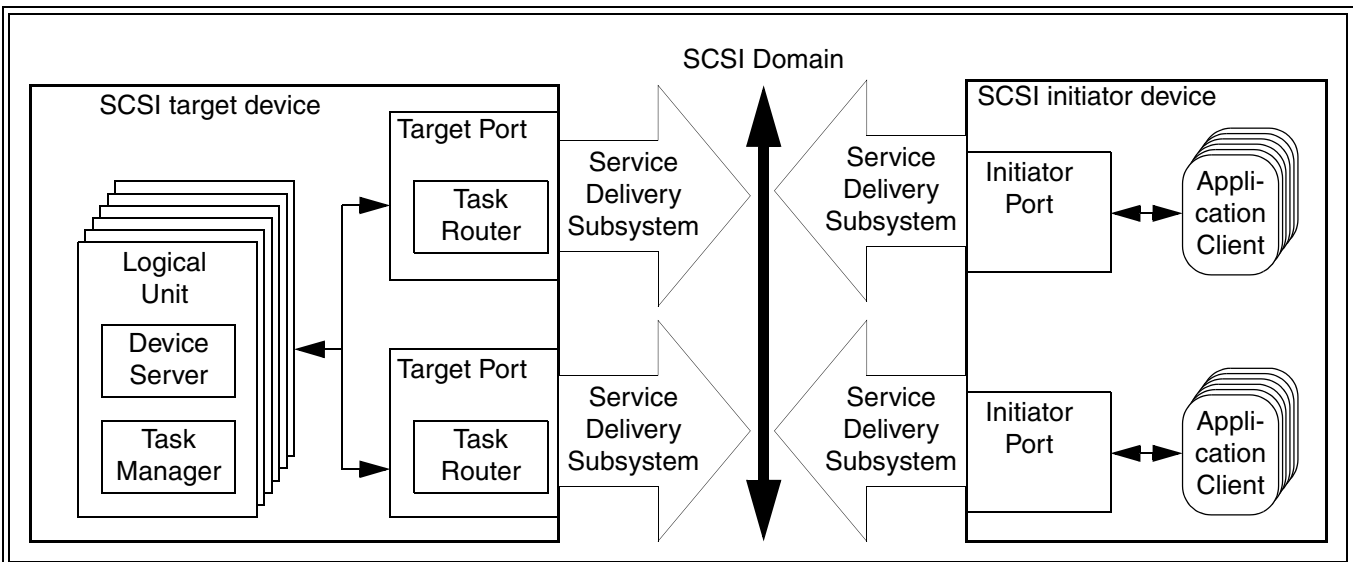


Figure 23 — SCSI target device and SCSI initiator device configured in a single SCSI domain

This model for application client determination of multiple SCSI target port configurations relies on information that is available only to the application clients via commands. The SCSI initiator ports in the SCSI initiator devices (figure 21) or SCSI initiator device (figure 22 and figure 23) are unable to distinguish the multiple SCSI target ports from individual SCSI target ports in two separate SCSI target devices.

4.13.7 SCSI target device view of a multiple port SCSI initiator device

This standard does not require a SCSI target device to have the ability to detect the presence of a SCSI initiator device with multiple SCSI initiator ports. Therefore, a SCSI target device handles a SCSI initiator device with multiple SCSI initiator ports exactly as it would handle multiple separate SCSI initiator devices (e.g., a SCSI target device handles the configurations shown in figure 22 and figure 23 in exactly the same way it handles the configuration shown in figure 21).

NOTE 4 - The implications of this view of a SCSI initiator device are more far reaching than are immediately apparent (e.g., after a SCSI initiator device makes an exclusive access reservation via one SCSI initiator port, access is denied to the other SCSI initiator port(s) on that same SCSI initiator device).

4.14 Model for dependent logical units

Optionally, the model for a logical unit (see 4.8) may include one or more unique logical units embedded within another logical unit. The embedded logical units are called dependent logical units (see 3.1.23). In such cases, the model hierarchy diagram in 4.8 is modified to become the diagram shown in figure 24.

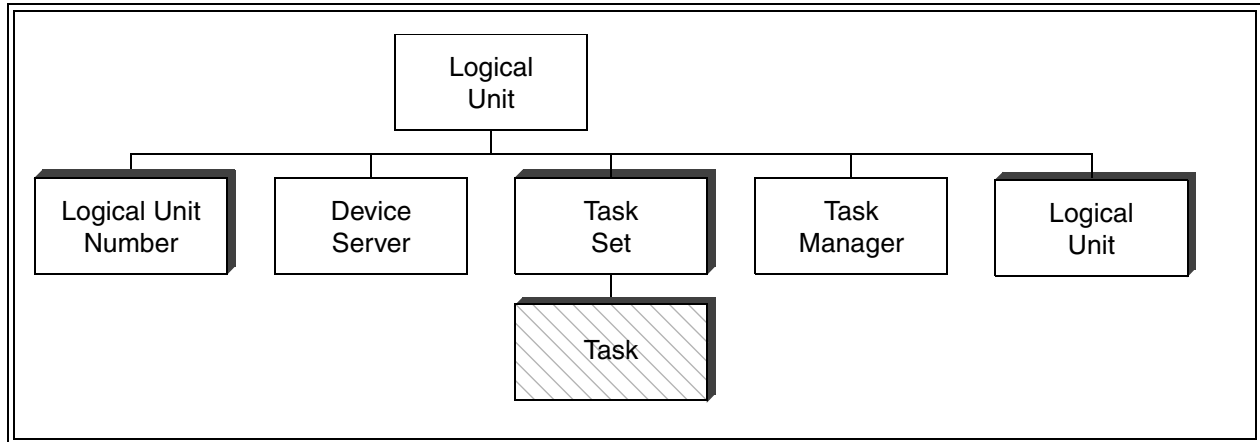


Figure 24 — Dependent logical unit model

When the dependent logical unit model is utilized, the hierarchical logical unit structure defined in 4.9.4 shall be used. If any logical unit within a SCSI target device includes dependent logical units, all logical unit numbers within the SCSI target device shall have the format described in 4.9.4. A device server that implements the hierarchical structure for dependent logical units described in this subclause shall set the HiSUP bit to one in the standard INQUIRY data returned by the logical unit with the logical unit number zero (see SPC-3).

The hierarchical logical unit structure is an inverted tree containing up to four addressable levels. The example in figure 25 is a three-level system that consists of:

- a) One SCSI initiator device that has three SCSI target devices attached in a single SCSI domain that is unable to add more SCSI target devices. One of the SCSI target devices is the level 1 SDDL (SDDL) with dependent logical units (SDDL). The level 1 SDDL has two SCSI target ports, one in each of the SCSI domains containing a SCSI initiator device;
- b) One SCSI initiator device has three SCSI target devices attached in a single SCSI domain that is able to add more SCSI target devices. One of the SCSI target devices is the level 1 SDDL; and
- c) The level 1 SDDL has three SCSI domains (called buses for hierarchical addressing purposes) with SCSI target devices attached and is capable of connecting more SCSI domains;
 - A) Two of the SCSI domains contain two SCSI target devices each and these SCSI domains are unable to add more SCSI target devices. One of the SCSI target devices is the level 2 SDDL;
 - B) One of the SCSI domains contains two SCSI target devices and is able to add more SCSI target devices; and
 - C) The level 2 SDDL has three SCSI domains with SCSI target devices attached and is capable of connecting more SCSI domains;
 - a) Two of the SCSI domains contain two SCSI target devices each and these SCSI domains are unable to add more SCSI target devices; and
 - b) One of the SCSI domains contains two SCSI target devices and is able to add more SCSI target devices.

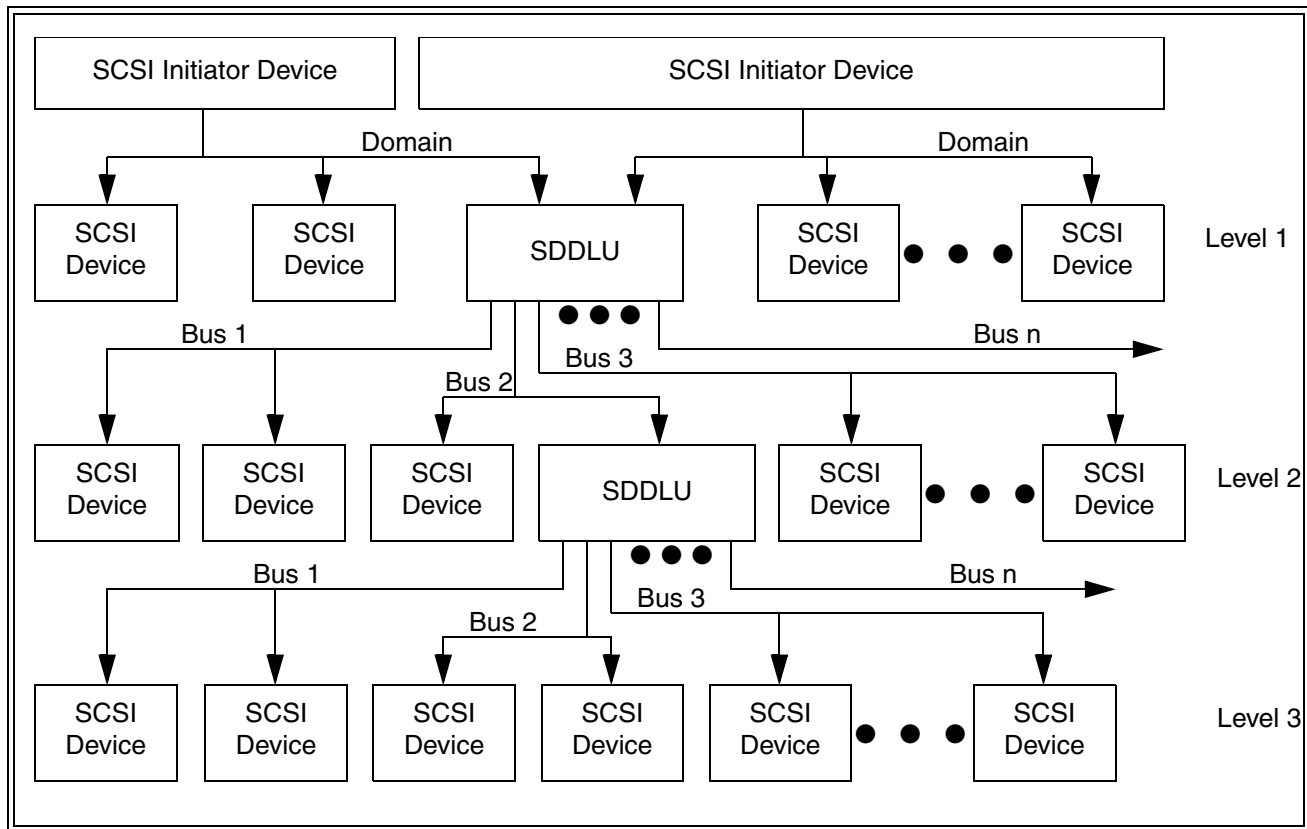


Figure 25 — Example of hierarchical system diagram

SCSI devices at each level in the tree are referenced by one of the following address methods:

- a) Logical unit address method (see 4.9.5);
- b) Peripheral device address method (see 4.9.6);
- c) Flat space addressing method (see 4.9.7); or
- d) Extended logical unit addressing method (see 4.9.8).

All peripheral device addresses, except LUN 0 (see 4.9.2), default to vendor specific values. All addressable entities may default to vendor specific values or may be defined by an application client (e.g., by the use of SCC-2 configuration commands).

Within the hierarchical system there may be SCSI target devices that have multiple logical units connected to them through separate SCSI initiator ports. The SCSI domains accessed by these SCSI initiator ports are referred to as buses. A SCSI target device that has SCSI devices attached to these buses shall assign numbers, other than zero, to those buses. The bus numbers shall be used as components of the logical unit numbers to the logical units attached to those buses, as described in 4.9.5 and 4.9.6.

SCSI target devices shall assign a bus number of zero to all the logical units under control by the SCSI target device that are not connected through a SCSI domain where the SCSI target device functions as a SCSI initiator device.

4.15 The SCSI model for distributed communications

The SCSI model for communications between distributed objects is based on the technique of layering as shown in figure 25.

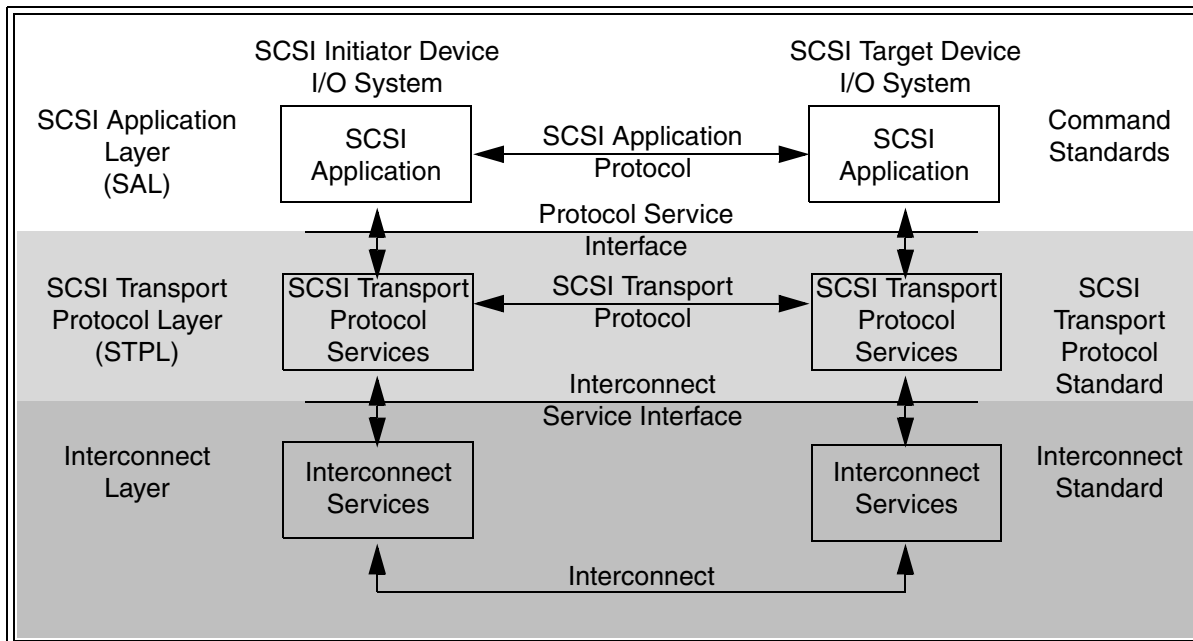


Figure 25 — Protocol service reference model

The layers comprising this model and the specifications defining the functionality of each layer are denoted by horizontal sequences. A layer consists of peer entities that communicate with one another by means of a protocol. Except for the interconnect layer, such communication is accomplished by invoking services provided by the adjacent layer. The following layers are defined.

SCSI application layer (SAL): Contains the clients and servers that originate and process SCSI I/O operations by means of a SCSI application protocol.

SCSI transport protocol layer (STPL): Consists of the services and protocols through which clients and servers communicate.

Interconnect layer: Comprised of the services, signaling mechanism and interconnect subsystem used for the physical transfer of data from sender to receiver. In the SCSI model, the interconnect layer is known as the service delivery subsystem.

The set of SCSI transport protocol services implemented by the service delivery subsystem identify external behavioral requirements that apply to SCSI transport protocol standards. While these SCSI transport protocol services may serve as a guide for designing reusable software or firmware that is adaptable to different SCSI transport protocols, there is no requirement for an implementation to provide the service interfaces specified in this standard.

The SCSI transport protocol service interface is defined in this standard in representational terms using SCSI transport protocol services. The SCSI transport protocol service interface implementation is defined in each SCSI transport protocol standard. The interconnect service interface is described as appropriate in each SCSI transport protocol standard.

Interactions between the SAL and STPL are defined with respect to the SAL and may originate in either layer. An outgoing interaction is modeled as a procedure call invoking an STPL service. An incoming interaction is modeled as a procedure call invoked by the STPL.

All procedure calls may be accompanied by parameters or data. Both types of interaction are described using the notation for procedures specified in 3.6.2. In this standard, input arguments are defined relative to the layer receiving an interaction (i.e., an input is a argument supplied to the receiving layer by the layer initiating the interaction).

The following types of service interactions between layers are defined.

SCSI transport protocol service request: A procedure call from the SAL invoking a service provided by the STPL.

SCSI transport protocol service indication: A procedure call from the STPL informing the SAL that an asynchronous event has occurred (e.g., the receipt of a peer-to-peer protocol transaction).

SCSI transport protocol service response: A procedure call to the STPL invoked by the SAL in response to a SCSI transport protocol service indication. A SCSI transport protocol service response may be invoked to return a reply from the invoking SAL to the peer SAL.

SCSI transport protocol service confirmation: A procedure call from the STPL notifying the SAL that a SCSI transport protocol service request has completed, has been terminated, or has failed to transit the interconnect layer. A confirmation may communicate parameters that indicate the completion status of the SCSI transport protocol service request or any other status. A SCSI transport protocol service confirmation may be used to convey a response from the peer SAL.

The services provided by an STPL are either confirmed or unconfirmed. A SAL service request invoking a confirmed service always results in a confirmation from the STPL.

Figure 26 shows the relationships between the four SCSI transport protocol service types.

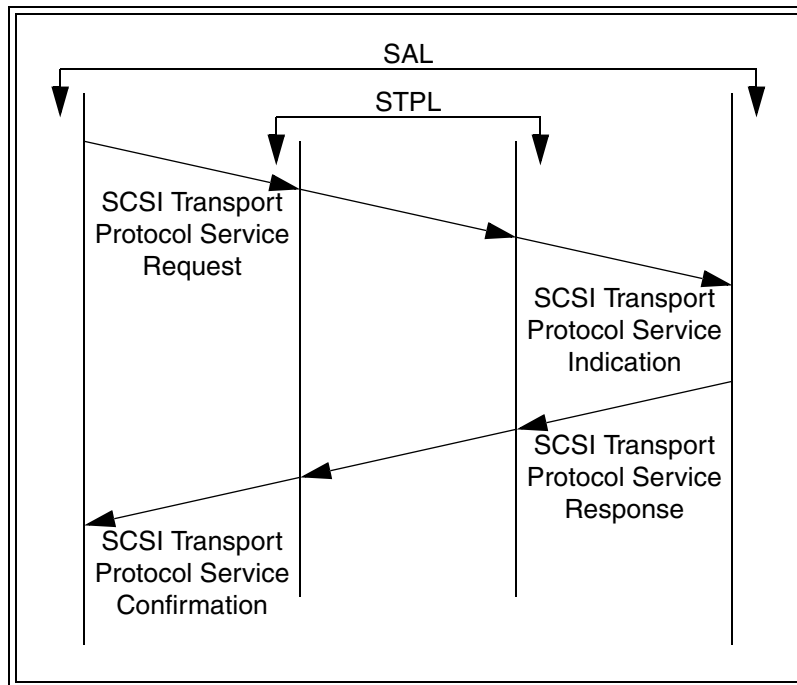


Figure 26 — SCSI transport protocol service model

Figure 27 shows how SCSI transport protocol services may be used to process a client-server request-response transaction at the SCSI application layer.

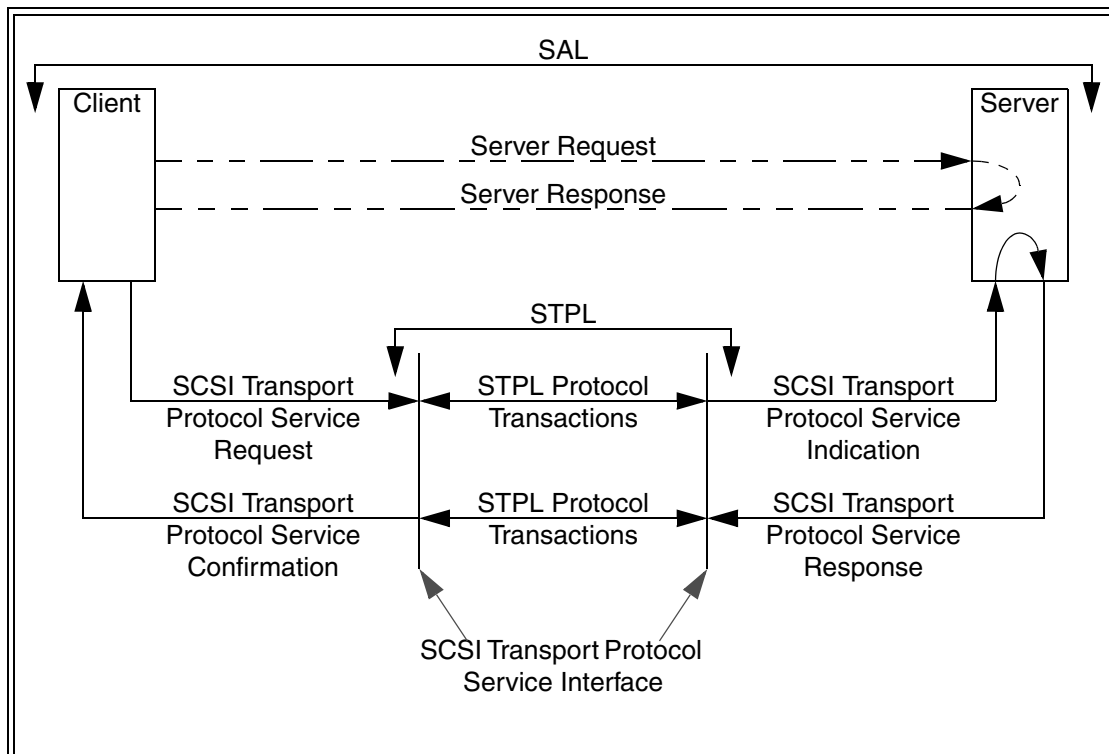


Figure 27 — Request-Response SAL transaction and related STPL services

The dashed lines in figure 27 show a SCSI application protocol transaction as it may appear to sending and receiving entities within the client and server. The solid lines in figure 27 show the corresponding SCSI transport protocol services and STPL transactions that are used to transport the data.

When a device server invokes a data transfer SCSI transport protocol service, the interactions required to transfer the data do not involve the application client. Only the STPL in the SCSI device that also contains the application client is involved. Figure 28 shows the relationships between the SCSI transport protocol service types involved in a data transfer request.

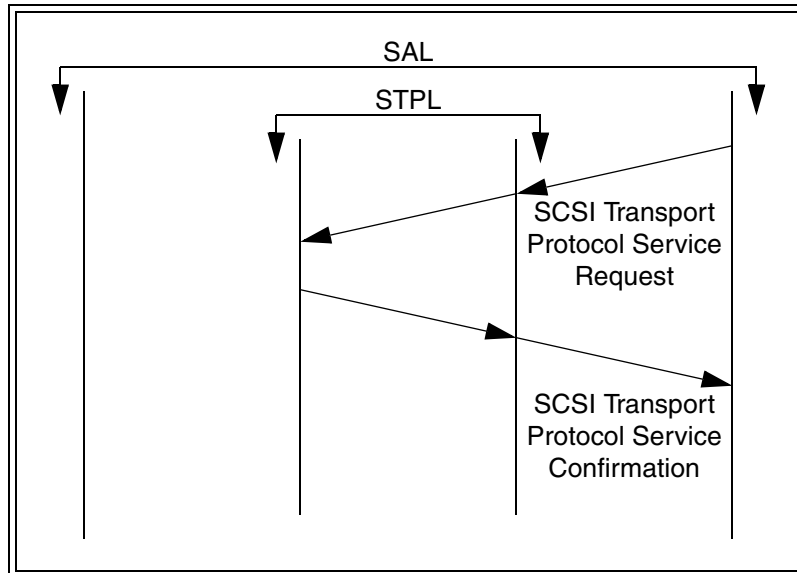


Figure 28 — SCSI transport protocol service model for data transfers

Figure 29 shows how SCSI transport protocol services may be used to process a device server data transfer transaction.

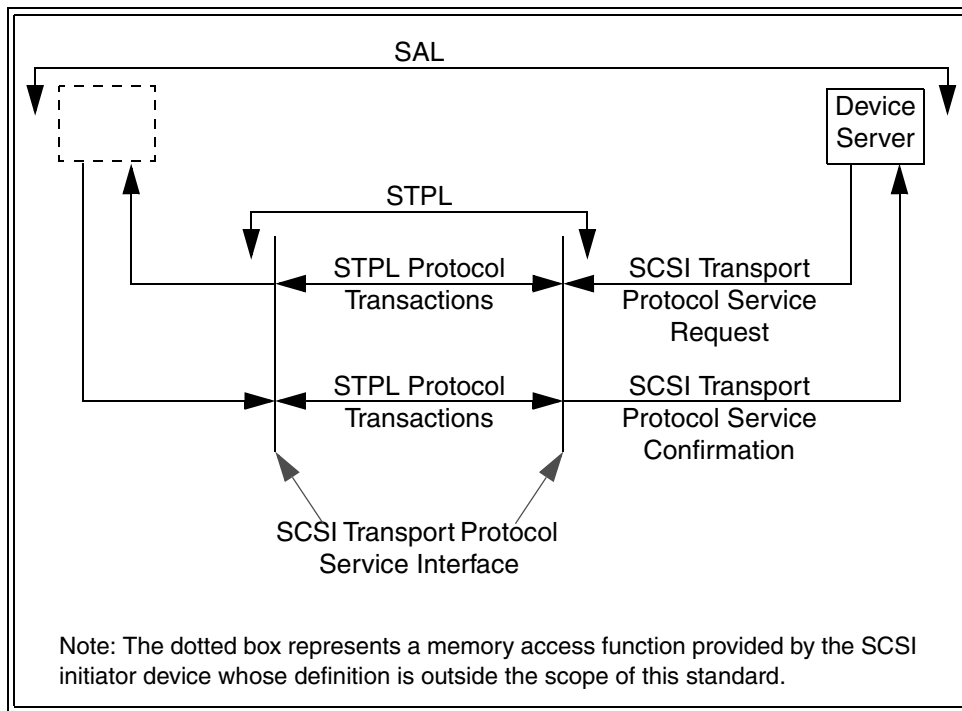


Figure 29 — Device server data transfer transaction and related STPL services

When a device server invokes a Terminate Data Transfer SCSI transport protocol service, the interactions required to complete the service do not involve the SCSI Transport Protocol Service Interface or the application client. Only the STPL in the SCSI device that also contains the device server is involved. Figure 30 shows the relationships between the SCSI transport protocol service types involved in a Terminate Data Transfer request.

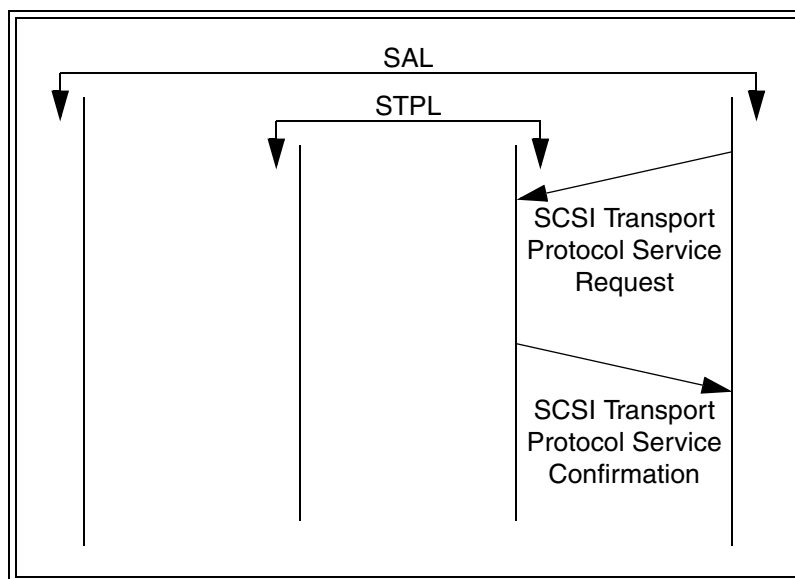


Figure 30 — SCSI transport protocol service model for Terminate Data Transfer

Figure 31 shows how SCSI transport protocol services may be used to process a device server Terminate Data Transfer transaction.

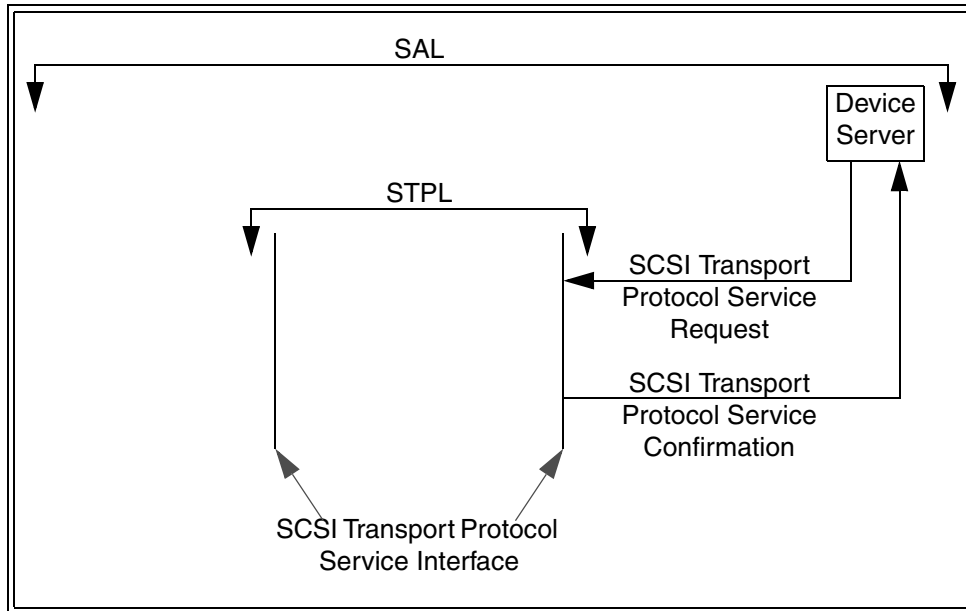


Figure 31 — Device server Terminate Data Transfer transaction and related STPL services

5 SCSI command model

5.1 The Execute Command procedure call

An application client requests the processing of a command by invoking the SCSI transport protocol services described in 5.4, the collective operation of which is modeled in the following procedure call:

Service Response = Execute Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size], [Data-Out Buffer], [Data-Out Buffer Size], [Command Reference Number], [Task Priority]), OUT ([Data-In Buffer], [Sense Data], [Sense Data Length], Status))

Input Arguments:

I_T_L_Q Nexus: The I_T_L_Q nexus identifying the task (see 4.12).

CDB: Command descriptor block (see 5.2).

Task Attribute: A value specifying one of the task attributes defined in 8.6. SCSI transport protocols may or may not provide the ability to specify a different task attribute for each task (see 8.6.1). For a task that processes linked commands, the Task Attribute shall be the value specified for the first command in a series of linked commands. The Task Attribute specified for the second and subsequent commands shall be ignored.

Data-In Buffer Size: The number of bytes available for data transfers to the Data-In Buffer (see 5.4.3).

Data-Out Buffer: A buffer containing command specific information to be sent to the logical unit (e.g., data or parameter lists needed to process the command). The buffer size is indicated by the Data-Out Buffer Size argument. The content of the Data-Out Buffer shall not change during the lifetime of the command (see 5.5) as viewed by the application client.

Data-Out Buffer Size: The number of bytes available for data transfers from the Data-Out Buffer (see 5.4.3).

Command Reference Number (CRN): When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one. The CRN shall be set to one for each I_T_L nexus involving the SCSI port after the SCSI port receives a hard reset or detects I_T nexus loss. The CRN shall be set to one after it reaches the maximum CRN value supported by the protocol. The CRN value zero shall be reserved for use as defined by the SCSI transport protocol. It is not an error for the application client to provide this argument when CRN is not supported by the SCSI transport protocol or logical unit.

Task Priority: The priority assigned to the task (see 8.7).

Output Arguments:

- Data-In Buffer:** A buffer to contain command specific information returned by the logical unit by the time of command completion. The **Execute Command** procedure call shall not return a status of GOOD, CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET unless the buffer contents are valid. The application client shall not assume that the buffer contents are valid unless the command completes with a status of GOOD, CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET. While some valid data may be present for other values of status, the application client should rely on additional information from the logical unit (e.g., sense data, to determine the state of the buffer contents). If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.
- Sense Data:** A buffer containing sense data returned in the same I_T_L_Q nexus transaction (see 3.1.46) as a CHECK CONDITION status (see 5.9.6). The buffer length is indicated by the Sense Data Length argument. If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.
- Sense Data Length:** The length in bytes of the Sense Data.
- Status:** A one-byte field containing command completion status (see 5.3). If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be undefined.

Service Response assumes one of the following values:

- TASK COMPLETE:** A logical unit response indicating that the task has ended. The Status argument shall have one of the values specified in 5.3 other than INTERMEDIATE or INTERMEDIATE-CONDITION MET.
- LINKED COMMAND COMPLETE:** Logical unit responses indicating that the task has not ended and that a linked command has completed successfully. As specified in 5.3, the Status argument shall have a value of INTERMEDIATE or INTERMEDIATE-CONDITION MET.
- SERVICE DELIVERY OR TARGET FAILURE:** The command has been ended due to a service delivery failure (see 3.1.113) or SCSI target device malfunction. All output parameters are invalid.

The SCSI transport protocol events corresponding to a response of TASK COMPLETE, LINKED COMMAND COMPLETE or SERVICE DELIVERY OR TARGET FAILURE shall be specified in each SCSI transport protocol standard.

An application client requests processing of a linked command by setting the LINK bit to one in the CDB CONTROL byte as specified in 5.2. The task attribute and task priority are determined by the Task Attribute argument and the Task Priority argument, respectively, specified for the first command in the series of linked commands. Upon receiving a response of LINKED COMMAND COMPLETE, an application client may issue the next command in the series through an **Execute Command** procedure call having the same I_T_L_Q nexus and omitting the Task Attribute argument. If the logical unit receives the next command in a series of linked commands before completing the current command in that linked command series, the overlapped command condition described in 5.9.3 shall result.

5.2 Command descriptor block (CDB)

The CDB defines the operation to be performed by the device server.

For all commands, if the logical unit detects an invalid parameter in the CDB, then the logical unit shall complete the command without altering the media information.

All CDBs shall have an OPERATION CODE as the first byte.

Some operation codes provide for modification of their operation based on a service action. In such cases, the combination of operation code value and service action code value may be modeled as a single, unique command determinate. The location of the SERVICE ACTION field in the CDB varies depending on the operation code value.

All CDBs shall contain a CONTROL byte (see table 20). The location of the CONTROL byte within a CDB depends on the CDB format (see SPC-3).

Table 20 — CONTROL byte

Bit	7	6	5	4	3	2	1	0
	Vendor specific		Reserved			NACA	Obsolete	LINK

All SCSI transport protocol standards shall define as mandatory the functionality needed for a logical unit to implement the NACA bit and LINK bit.

The NACA (Normal ACA) bit specifies whether an auto contingent allegiance (ACA) is established if the command returns with CHECK CONDITION status. An NACA bit set to one specifies that an ACA shall be established. An NACA bit set to zero specifies that an ACA shall not be established. The actions for ACA are specified in 5.9.2. Actions that may be required when an ACA is not established are described in 5.9.1. All logical units shall implement support for the NACA value of zero and may support the NACA value of one (i.e., ACA). The ability to support a NACA value of one is indicated with the NORMACA bit in the standard INQUIRY data (see SPC-3).

If the NACA bit is set to one but the logical unit does not support ACA, the logical unit shall complete the command with a CHECK CONDITION status, sense key of ILLEGAL REQUEST, an additional sense code of INVALID FIELD IN CDB.

The LINK bit is used to continue the task across multiple commands. Support for the LINK bit is optional. The application client sets the LINK bit to one to specify a request for continuation of the task across two or more commands. If the LINK bit is set to one and the command completes successfully, a logical unit that supports the LINK bit shall continue the task and return a status of INTERMEDIATE or INTERMEDIATE-CONDITION MET and a service response of LINKED COMMAND COMPLETE (see 5.3). The logical unit shall complete the command with a status of CHECK CONDITION and a sense key of ILLEGAL REQUEST if the LINK bit is set to one and the logical unit does not support linked commands.

Bit 1 provides an obsolete way to request interrupts between linked commands.

5.3 Status

5.3.1 Status codes

The status codes are specified in table 21. Status shall be sent from the device server to the application client whenever a command ends with a service response of TASK COMPLETE or LINKED COMMAND COMPLETE.

Table 21 — Status codes

Status Code	Status	Task Ended	Service Response
00h	GOOD	Yes	TASK COMPLETE
02h	CHECK CONDITION	Yes	TASK COMPLETE
04h	CONDITION MET	Yes	TASK COMPLETE
08h	BUSY	Yes	TASK COMPLETE
10h	INTERMEDIATE	No	LINKED COMMAND COMPLETE
14h	INTERMEDIATE-CONDITION MET	No	LINKED COMMAND COMPLETE
18h	RESERVATION CONFLICT	Yes	TASK COMPLETE
22h	Obsolete		
28h	TASK SET FULL	Yes	TASK COMPLETE
30h	ACA ACTIVE	Yes	TASK COMPLETE
40h	TASK ABORTED	Yes	TASK COMPLETE
All other codes	Reserved		

Definitions for each status code are as follows:

GOOD. This status indicates that the device server has successfully completed the task.

CHECK CONDITION. This status indicates that sense data has been delivered in the buffer defined by the Sense Data argument to the **Execute Command** procedure call (see 5.9.6). Additional actions that are required when CHECK CONDITION status is returned are described in 5.9.1.

CONDITION MET. This status shall be returned whenever the requested operation specified by an unlinked command is satisfied (see the PRE-FETCH commands in the SBC standard).

BUSY. This status indicates that the logical unit is busy. This status shall be returned whenever a logical unit is temporarily unable to accept a command. The recommended application client recovery action is to issue the command again at a later time. If the UA_INTLCK_CTRL field in the Control mode page contains 11b (see SPC-3), termination of a command with BUSY status shall cause an unit attention condition to be established for the SCSI initiator port that sent the command with an additional sense code of PREVIOUS BUSY STATUS unless a PREVIOUS BUSY STATUS unit attention condition already exists.

INTERMEDIATE. This status or INTERMEDIATE-CONDITION MET shall be returned for each successfully completed command in a series of linked commands (except the last command), unless the command is terminated with CHECK CONDITION, RESERVATION CONFLICT, TASK SET FULL, or BUSY status. If INTERMEDIATE or INTERMEDIATE-CONDITION MET status is not returned, the series of linked commands is terminated and the task is ended. This status is the equivalent of GOOD status for linked commands.

INTERMEDIATE-CONDITION MET. This status is returned whenever the requested operation specified by a linked command is satisfied (see the PRE-FETCH commands in the SBC standard), unless the command is terminated with CHECK CONDITION, RESERVATION CONFLICT, TASK SET FULL, or BUSY status. If INTERME-

DIATE or INTERMEDIATE-CONDITION MET status is not returned, the series of linked commands is terminated and the task is ended.

RESERVATION CONFLICT. This status shall be returned whenever a SCSI initiator port attempts to access a logical unit or an element of a logical unit in a way that conflicts with an existing reservation. (See the PERSISTENT RESERVE OUT command and PERSISTENT RESERVE IN command in SPC-3).

If the UA_INTLCK_CTRL field in the Control mode page contains 11b (see SPC-3), termination of a command with RESERVATION CONFLICT status shall cause an unit attention condition to be established for the SCSI initiator port that sent the command with an additional sense code of PREVIOUS RESERVATION CONFLICT STATUS unless a PREVIOUS RESERVATION CONFLICT STATUS unit attention condition already exists.

TASK SET FULL. This status shall be implemented by all logical units.

When the logical unit has at least one task in the task set for a SCSI initiator port and a lack of task set resources prevents accepting a received task from that SCSI initiator port into the task set, TASK SET FULL shall be returned. When the logical unit has no task in the task set for a SCSI initiator port and a lack of task set resources prevents accepting a received task from that SCSI initiator port into the task set, BUSY should be returned.

The logical unit should allow at least one command in the task set for each supported SCSI initiator port that has identified itself to the SCSI target port in a SCSI transport protocol specific manner or by the successful transmission of a command.

If the UA_INTLCK_CTRL field in the Control mode page contains 11b (see SPC-3), termination of a command with TASK SET FULL status shall cause an unit attention condition to be established for the SCSI initiator port that sent the command with an additional sense code of PREVIOUS TASK SET FULL STATUS unless a PREVIOUS TASK SET FULL STATUS unit attention condition already exists.

ACA ACTIVE. This status shall be returned when an ACA exists within a task set and a SCSI initiator port issues a command for that task set when at least one of the following is true:

- a) There is a task with the ACA attribute (see 8.6.5) in the task set;
- b) The SCSI initiator port issuing the command did not cause the ACA condition; or
- c) The task created to process the command did not have the ACA attribute and the NACA bit was set to one in the CDB CONTROL byte of the faulting command (see 5.9.1).

The SCSI initiator port may reissue the command after the ACA condition has been cleared.

TASK ABORTED. This status shall be returned when a task is aborted by another SCSI initiator port and the Control mode page TAS bit is set to one (see 5.7.3).

5.3.2 Status precedence

If a device server detects that more than one of the following conditions applies to a completed task, it shall select the condition to report based on the following precedence:

- 1) An ACA ACTIVE status;
- 2) A CHECK CONDITION status for any of the following unit attention conditions (i.e., with a sense key of UNIT ATTENTION and one of the following additional sense codes):
 - A) POWER ON, RESET, OR BUS DEVICE RESET OCCURRED;
 - B) POWER ON OCCURRED;
 - C) SCSI BUS RESET OCCURRED;
 - D) BUS DEVICE RESET FUNCTION OCCURRED;
 - E) DEVICE INTERNAL RESET; or
 - F) I_T NEXUS LOSS OCCURRED;
- 3) A RESERVATION CONFLICT status;
and
- 4) A status of:
 - A) CHECK CONDITION, for any reason not listed in 1);
 - B) GOOD;
 - C) CONDITION MET;
 - D) INTERMEDIATE;
 - E) INTERMEDIATE-CONDITION MET; or
 - F) TASK ABORTED.

NOTE 5 - The names of the unit attention conditions listed in the subclause (e.g., SCSI BUS RESET OCCURRED) are based on usage in previous versions of this standard. The use of these unit attention condition names is not to be interpreted as a description of how the unit attention conditions are represented by any given SCSI transport protocol.

A device server may report the following status codes with any level of precedence:

- a) BUSY status;
- b) TASK SET FULL status; or
- c) CHECK CONDITION status for a sense key of ILLEGAL REQUEST.

5.4 SCSI transport protocol services in support of Execute Command

5.4.1 Overview

The SCSI transport protocol services that support the **Execute Command** procedure call are described in 5.4. Two groups of SCSI transport protocol services are described. The SCSI transport protocol services that support the request and confirmation for the **Execute Command** procedure call are described in 5.4.2. The SCSI transport protocol services that support the data transfers associated with processing a command are described in 5.4.3.

5.4.2 Execute Command request/confirmation SCSI transport protocol services

All SCSI transport protocol standards shall define the SCSI transport protocol specific requirements for implementing the **Send SCSI Command** SCSI transport protocol service request and the **Command Complete Received** confirmation. Support for the **SCSI Command Received** indication and **Send Command Complete** response by a SCSI transport protocol standard is optional. All SCSI I/O systems shall implement these SCSI transport protocols as defined in the applicable SCSI transport protocol specification.

SCSI Transport Protocol Service Request:

**Send SCSI Command (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Data-In Buffer Size],
[Data-Out Buffer], [Data-Out Buffer Size], [Command Reference Number],
[Task Priority], [First Burst Enabled]))**

Input Arguments:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- CDB:** Command descriptor block (see 5.2).
- Task Attribute:** A value specifying one of the task attributes defined in 8.6. For specific requirements on the Task Attribute argument see 5.1.
- Data-In Buffer Size:** The number of bytes available for data transfers to the Data-In Buffer (see 5.4.3).
- Data-Out Buffer:** A buffer containing command specific information to be sent to the logical unit (e.g., data or parameter lists needed to process the command (see 5.1)). The content of the Data-Out Buffer shall not change during the lifetime of the command (see 5.5) as viewed by the application client.
- Data-Out Buffer Size:** The number of bytes available for data transfers from the Data-Out Buffer (see 5.4.3).
- Command Reference Number (CRN):** When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one (see 5.1).
- Task Priority:** The priority assigned to the task (see 8.7).
- First Burst Enabled:** An argument specifying that a SCSI transport protocol specific number of bytes from the Data-Out Buffer shall be delivered to the logical unit without waiting for the device server to invoke the **Receive Data-Out** SCSI transport protocol service.

SCSI Transport Protocol Service Indication:

SCSI Command Received (IN (I_T_L_Q Nexus, CDB, Task Attribute, [Command Reference Number], [Task Priority], [First Burst Enabled]))

Input Arguments:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- CDB:** Command descriptor block (see 5.2).
- Task Attribute:** A value specifying one of the task attributes defined in 8.6. For specific requirements on the Task Attribute argument see 5.1.
- Command Reference Number (CRN):** When this argument is used, all sequential commands of an I_T_L nexus shall include a CRN argument that is incremented by one (see 5.1).
- Task Priority:** The priority assigned to the task (see 8.7).
- First Burst Enabled:** An argument specifying that a SCSI transport protocol specific number of bytes from the Data-Out Buffer are being delivered to the logical unit without waiting for the device server to invoke the **Receive Data-Out** SCSI transport protocol service.

SCSI Transport Protocol Service Response (from device server):

Send Command Complete (IN (I_T_L_Q Nexus, [Sense Data], [Sense Data Length], Status, Service Response))

Input Arguments:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- Sense Data:** If present, this argument instructs the SCSI target port to return sense data to the SCSI initiator port (see 5.9.6).
- Sense Data Length:** The length in bytes of the sense data to be returned to the SCSI initiator port.
- Status:** Command completion status (see 5.1).
- Service Response:** Possible service response information for the command (see 5.1).

SCSI Transport Protocol Service Confirmation:

Command Complete Received (IN (I_T_L_Q Nexus, [Data-In Buffer], [Sense Data], [Sense Data Length], Status, Service Response))

Input Arguments:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- Data-In Buffer:** A buffer containing command specific information returned by the logical unit on command completion (see 5.1).
- Sense Data:** Sense data returned in the same I_T_L_Q nexus transaction (see 3.1.46) as a CHECK CONDITION status (see 5.9.6).
- Status:** Command completion status (see 5.1).
- Sense Data Length:** The length in bytes of the received sense data.
- Service Response:** Service response for the command (see 5.1).

5.4.3 Data transfer SCSI transport protocol services

5.4.3.1 Introduction

The data transfer services described in 5.4.3 provide mechanisms for moving data to and from the SCSI initiator port in response to commands transmitted using the **Execute Command** procedure call. All SCSI transport protocol standards shall define the protocols required to implement these services.

The application client's Data-In Buffer and/or Data-Out Buffer each appears to the device server as a single, logically contiguous block of memory large enough to hold all the data required by the command (see figure 32). This standard allows either unidirectional or bidirectional data transfer. The processing of a command may require the transfer of data from the application client using the Data-Out Buffer, or to the application client using the Data-In Buffer, or both to and from the application client using both the Data-In Buffer and the Data-Out Buffer.

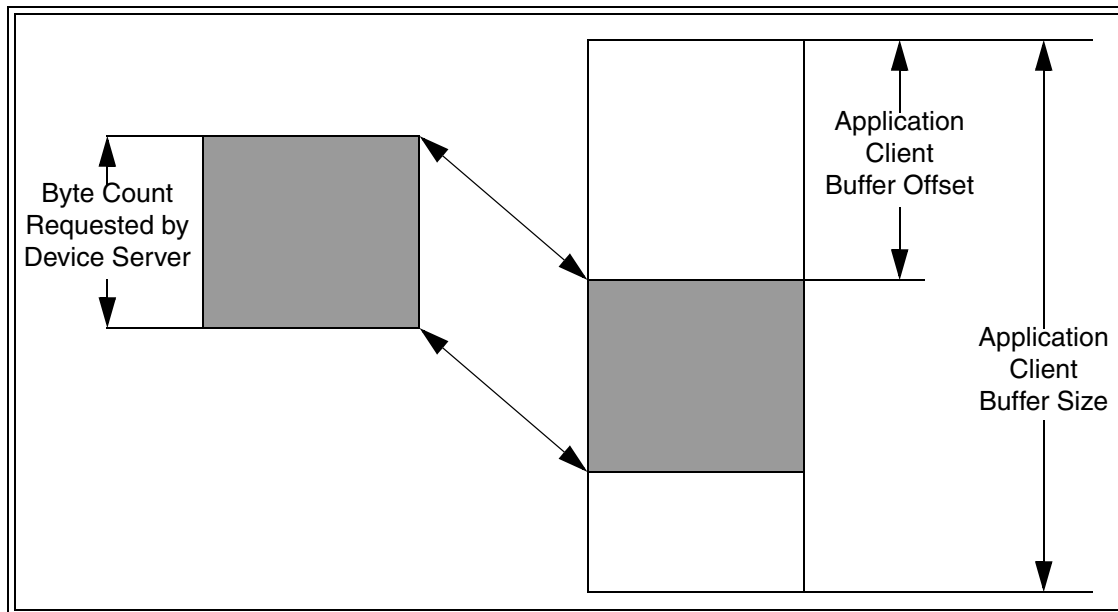


Figure 32 — Model for Data-In and Data-Out data transfers

This standard assumes that the buffering resources available to the logical unit are limited and may be less than the amount of data that is capable of being transferred in one command. Such data needs to be moved between the application client and the media in segments that are smaller than the transfer size specified in the command. The amount of data moved per segment is usually a function of the buffering resources available to the logical unit. Figure 32 shows the model for such incremental data transfers.

SCSI transport protocols may allow logical units to accept the initial portion of the Data-Out Buffer data, called the first burst, along with the command without waiting for the device server to invoke the **Receive Data-Out** SCSI transport protocol service. This is modeled using **Receive Data-Out** protocol service calls for which the SCSI transport protocol may have moved the first burst prior to the call.

SCSI transport protocols that define a first burst capability shall include the First Burst Enabled argument in their definitions for the **Send SCSI Command** and **SCSI Command Received** procedure calls. Logical units that implement the first burst capability shall implement the FIRST BURST SIZE field in the Disconnect-Reconnect mode page (see SPC-3).

The movement of data between the application client and device server is controlled by the following arguments:

- Application Client Buffer Size:** The total number of bytes in the application client's buffer (Data-In or Data-Out).
- Application Client Buffer Offset:** Offset in bytes from the beginning of the application client's buffer (Data-In or Data-Out) to the first byte of transferred data.
- Byte Count Requested by Device Server:** Number of bytes to be moved by the data transfer request.

For any specific data transfer SCSI transport protocol service request, the **Byte Count Requested by Device Server** is less than or equal to the combination of **Application Client Buffer Size** minus the **Application Client Buffer Offset**.

If a SCSI transport protocol supports random buffer access, the offset and byte count specified for each data segment to be transferred may overlap. In this case the total number of bytes moved for a command is not a reliable indicator of highest byte transferred and shall not be used by a SCSI initiator device or SCSI target device implementation to determine whether all data has been transferred.

All SCSI transport protocol standards shall define support for a resolution of one byte for the above arguments. A SCSI initiator device shall support a resolution of one byte. A SCSI target device may support any resolution.

Editor's Note 1 - ROW: Concerns have been raised about the requirements in the above paragraph being overly restrictive. Depending on one's point of view, several of the packetized protocols have restrictions on data alignment with a resolution greater than one byte. The restrictions on buffer length allow even odd byte granularity.

It is possible that the above restrictions have been placed at a layer of the model that is no longer appropriate. This layer may be viewed as applying to the internal mechanics of a HBA, whereas the intent is that user data buffers identified in an **Execute Command** procedure call be allowed have one byte alignment and length granularity. The layers that interface between the **Execute Command** procedure call would/should be obliged to accommodate differences between user data buffer alignments and the potentially more restrictive requirements commonly applied at this layer.

Random buffer access occurs when the device server requests data transfers to or from segments of the application client's buffer that have an arbitrary offset and byte count. Buffer access is sequential when successive transfers access a series of monotonically increasing, adjoining buffer segments. Support for random buffer access by a SCSI transport protocol standard is optional. A device server implementation designed for any SCSI transport protocol implementation should be prepared to use sequential buffer access when necessary.

The STPL confirmed services specified in 5.4.3.2 and 5.4.3.3 are used by the device server to request the transfer of command data to or from the application client. The SCSI initiator device SCSI transport protocol service interactions are unspecified.

This standard provides only for the transfer phases to be sequential. Provision for overlapping transfer phases is outside the scope of this standard.

5.4.3.2 Data-In delivery service

Request:

Send Data-In (IN (I_T_L_Q Nexus, Device Server Buffer, Application Client Buffer Offset, Request Byte Count))

Argument descriptions:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- Device Server Buffer:** The buffer in the device server from which data is to be transferred.
- Application Client Buffer Offset:** Offset in bytes from the beginning of the application client's buffer to the first byte of transferred data.
- Request Byte Count:** Number of bytes to be moved by this request.

Confirmation:

Data-In Delivered (IN (I_T_L_Q Nexus, Delivery Result))

This confirmation notifies the device server that the specified data was successfully delivered to the application client buffer, or that a service delivery subsystem error occurred while attempting to deliver the data.

Argument descriptions:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- Delivery Result:** an encoded value representing one of the following:
 - DELIVERY SUCCESSFUL: The data was delivered successfully.
 - DELIVERY FAILURE: A service delivery subsystem error occurred while attempting to deliver the data.

5.4.3.3 Data-Out delivery service

Request:

Receive Data-Out (IN (I_T_L_Q Nexus, Application Client Buffer Offset, Request Byte Count, Device Server Buffer))

Argument descriptions:

- I_T_L_Q Nexus:** The I_T_L_Q nexus identifying the task (see 4.12).
- Device Server Buffer:** The buffer in the device server to which data is to be transferred.
- Application Client Buffer Offset:** Offset in bytes from the beginning of the application client's buffer to the first byte of transferred data.
- Request Byte Count:** Number of bytes to be moved by this request.

If the **SCSI Command Received** SCSI transport protocol service included a First Burst Enabled argument and random buffer access is not supported, first burst data shall be transferred to the Device Server Buffer until all first burst data has been transferred. If the **SCSI Command Received** SCSI transport protocol service included a First Burst Enabled argument and random buffer access is supported, first burst data should be transferred to the Device Server Buffer but first burst data may be re-transferred across the service delivery subsystem.

Confirmation:**Data-Out Received (IN (I_T_L_Q Nexus, Delivery Result))**

This confirmation notifies the device server that the requested data has been successfully delivered to its buffer, or that a service delivery subsystem error occurred while attempting to receive the data.

Argument descriptions:

I_T_L_Q Nexus: The I_T_L_Q nexus identifying the task (see 4.12).

Delivery Result: an encoded value representing one of the following:

DELIVERY SUCCESSFUL: The data was delivered successfully.

DELIVERY FAILURE: A service delivery subsystem error occurred while attempting to receive the data.

5.4.3.4 Terminate Data Transfer service

The **Terminate Data Transfer** SCSI transport protocol service allows a device server to specify that one or more **Send Data-In** or **Receive Data-Out** SCSI transport protocol service requests be terminated by a SCSI target port.

Request:**Terminate Data Transfer (IN (Nexus))**

Argument description:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

The SCSI target port terminates all transfer service requests for the specified nexus (e.g., if an I_T_L nexus is specified, then the SCSI target port terminates all transfer service requests from the logical unit for the specified SCSI initiator port).

Confirmation:**Data Transfer Terminated (IN (Nexus))**

Argument description:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

This confirmation is returned in response to a **Terminate Data Transfer** request whether or not the specified nexus existed in the SCSI target port when the request was received. After a **Data Transfer Terminated** SCSI transport protocol service confirmation has been sent in response to a **Terminate Data Transfer** SCSI transport protocol service request, **Data-In Delivered** or **Data-Out Received** SCSI transport protocol service confirmations shall not be sent for the tasks specified by the nexus.

5.5 Task and command lifetimes

This subclause specifies the events delimiting the beginning and end (i.e., lifetime) of a task or pending command from the viewpoint of the device server and application client.

The device server shall create a task upon receiving a **SCSI Command Received** indication unless the command represents a continuation of a linked command as described in 5.1.

The task shall exist until:

- a) The device server sends a SCSI transport protocol service response for the task of TASK COMPLETE; or
- b) The task is aborted as described in 5.7.

The application client assumes that the task exists and maintains an application client task to interact with the task from the time the **Send SCSI Command** SCSI transport protocol service request is invoked until it receives one of the following SCSI target device responses:

- a) A service response of TASK COMPLETE for that task;
- b) Notification of an unit attention condition with one of the following additional sense codes;
 - A) COMMANDS CLEARED BY ANOTHER INITIATOR (if in reference to the task set containing the task);
 - B) Any additional sense code whose ADDITIONAL SENSE CODE field contains 29h (e.g., POWER ON, RESET, OR BUS DEVICE RESET OCCURRED; POWER ON OCCURRED; SCSI BUS RESET OCCURRED; BUS DEVICE RESET FUNCTION OCCURRED; DEVICE INTERNAL RESET; or I_T NEXUS LOSS OCCURRED);
- c) A service response of SERVICE DELIVERY OR TARGET FAILURE for the command. The task associated with the failed command shall be ended;
- d) A service response of FUNCTION COMPLETE following an ABORT TASK task management function directed to the specified task;
- e) A service response of FUNCTION COMPLETE following an ABORT TASK SET or a CLEAR TASK SET task management function directed to the task set containing the specified task; or
- f) A service response of FUNCTION COMPLETE in response to a LOGICAL UNIT RESET task management function directed to the logical unit.

NOTE 6 - The names of the unit attention conditions listed in the subclause (e.g., SCSI BUS RESET OCCURRED) are based on usage in previous versions of this standard. The use of these unit attention condition names is not to be interpreted as a description of how the unit attention conditions are represented by any given SCSI transport protocol.

To the application client, the command is pending from the time it calls the **Send SCSI Command** SCSI transport protocol service until one of the above responses or a service response of LINKED COMMAND COMPLETE is received.

When a SCSI transport protocol does not require state synchronization (see 4.6.2), there may be a time skew between the completion of a device server request-response transaction as seen by the application client and device server. As a result, the lifetime of a task or command as it appears to the application client normally is different from the lifetime observed by the device server.

Some commands (e.g., commands with immediate bits like SEND DIAGNOSTIC, or write commands when a write cache is enabled) start background operations that operate after the task containing the command is no longer in the task set. These operations may be aborted by power on, hard resets, or logical unit resets. These operations shall not be aborted by I_T nexus loss.

Background operations may generate deferred errors that are reported in the sense data for a subsequent completed command (see SPC-3). A deferred error may be cleared by power on, hard reset, logical unit reset, or by I_T nexus loss involving the SCSI initiator port from which the command that spawned the background

operation was received, if the device server has retained that information. Deferred errors should not be cleared by I_T nexus loss.

5.6 Task management function lifetime

The application client assumes that the task management function is in process from the time the **Send Task Management Request** SCSI transport protocol service request is invoked until it receives one of the following SCSI target device responses:

- a) A service response of FUNCTION COMPLETE, FUNCTION SUCCEEDED, FUNCTION REJECTED, or SERVICE DELIVERY OR TARGET FAILURE is received for that task management function; or
- b) Notification of an unit attention condition with any additional sense code whose additional sense code field contains 29h (e.g., POWER ON, RESET, OR BUS DEVICE RESET OCCURRED; POWER ON OCCURRED; SCSI BUS RESET OCCURRED; BUS DEVICE RESET FUNCTION OCCURRED; DEVICE INTERNAL RESET).

NOTE 7 - The names of the unit attention conditions listed in the subclause (e.g., SCSI BUS RESET OCCURRED) are based on usage in previous versions of this standard. The use of these unit attention condition names is not to be interpreted as a description of how the unit attention conditions are represented by any given SCSI transport protocol.

5.7 Aborting tasks

5.7.1 Mechanisms that cause tasks to be aborted

A task is aborted when an event or SCSI initiator device action causes termination of the task prior to its successful completion.

The following events cause a task or several tasks to be aborted:

- a) The return of an **Execute Command** service response of SERVICE DELIVERY OR TARGET FAILURE as described in 5.1;
- b) A logical unit reset (see 6.3.3);
- c) A hard reset (see 6.3.2);
- d) A power on condition; or
- e) SCSI transport protocol specific events.

An action transmitted via a SCSI initiator port may abort task(s) created via the SCSI initiator port itself, task(s) created via another SCSI initiator port, or both its own tasks and tasks created via another SCSI initiator port.

The following actions affect only the task(s) created via the SCSI initiator port that transmits the action:

- a) Completion of an ABORT TASK task management function directed to the specified task;
- b) Completion of an ABORT TASK SET task management function under the conditions specified in 7.3; or
- c) A command completes with a CHECK CONDITION status without establishing an ACA condition (see 5.9.1.3) or establishing an ACA condition (see 5.9.2.1) while the Control mode page (see SPC-3) contains fields that are set as follows:
 - A) The QERR field set to 01b and the TST field set to 001b; or
 - B) The QERR field set to 11b.

The following actions affect the task(s) created via the SCSI initiator port that transmits the action and/or task(s) created via other SCSI initiator ports:

- a) Completion of a CLEAR TASK SET task management function referencing the task set containing the specified task;
- b) A command completes with a CHECK CONDITION status with or without establishing an ACA condition and the QERR field was set to 01b and the TST field was set to 000b in the Control mode page (see SPC-3);
- c) Completion of a PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action with a reservation key that is associated with the SCSI initiator port that created the task (see SPC-3); or
- d) Completion of a LOGICAL UNIT RESET task management function (see 7.6) directed to the logical unit.

If one or more tasks are cleared or aborted, the affected tasks are also cleared from the initiator ports in a manner that is outside the scope of this standard.

5.7.2 When a SCSI initiator port aborts its own tasks

When a SCSI initiator port causes its own task(s) to be aborted, no notification that the task(s) have been aborted shall be returned to the SCSI initiator port other than the completion response for the command or task management function action that caused the task(s) to be aborted and notification(s) associated with related effects of the action (e.g., a reset unit attention condition).

5.7.3 When a SCSI initiator port aborts tasks from other SCSI initiator ports

When a SCSI initiator port causes the task(s) of another SCSI initiator port to be aborted, the other SCSI initiator port shall be notified that the task(s) have been aborted. The method of notification shall depend on the setting of the TAS bit in the Control mode page (see SPC-3) that applies to the other SCSI initiator port.

If the TAS bit is set to zero, the method of notification shall be an unit attention condition. The additional sense code set for the unit attention condition depends on the action that caused the task(s) to be aborted.

If the TAS bit is set to one, the method of notification shall be the termination of each aborted task with a TASK ABORTED status. The COMMANDS CLEARED BY ANOTHER INITIATOR unit attention condition shall not be established, however, the establishment of any other applicable unit attention condition shall not be affected.

When a logical unit is aborting one or more tasks from a SCSI initiator port with the TASK ABORTED status it should complete all of those tasks before entering additional tasks from that SCSI initiator port into the task set.

5.8 Command processing examples

5.8.1 Unlinked command example

An unlinked command is used to show the events associated with the processing of a single device service request (see figure 33). This example does not include error or exception conditions.




Figure 33 — Command processing events

The numbers in figure 33 identify the events described as follows:

1. The application client task performs an **Execute Command** procedure call by invoking the **Send SCSI Command** SCSI transport protocol service to send the CDB and other input parameters to the logical unit.
2. The device server is notified through a **SCSI Command Received** indication containing the CDB and command parameters. A task is created and entered into the task set. The device server may invoke the appropriate data delivery service one or more times to complete command processing.
3. The task ends upon completion of the command. On command completion, the **Send Command Complete** SCSI transport protocol service is invoked to return a status of GOOD and a service response of TASK COMPLETE.
4. A confirmation of **Command Complete Received** is passed to the application client task by the SCSI initiator port.

5.8.2 Linked command example

A task may consist of multiple commands linked together. After the logical unit notifies the application client that a linked command has successfully completed, the application client issues the next command in the series.

The example in figure 34 shows the events in a series of two linked commands.

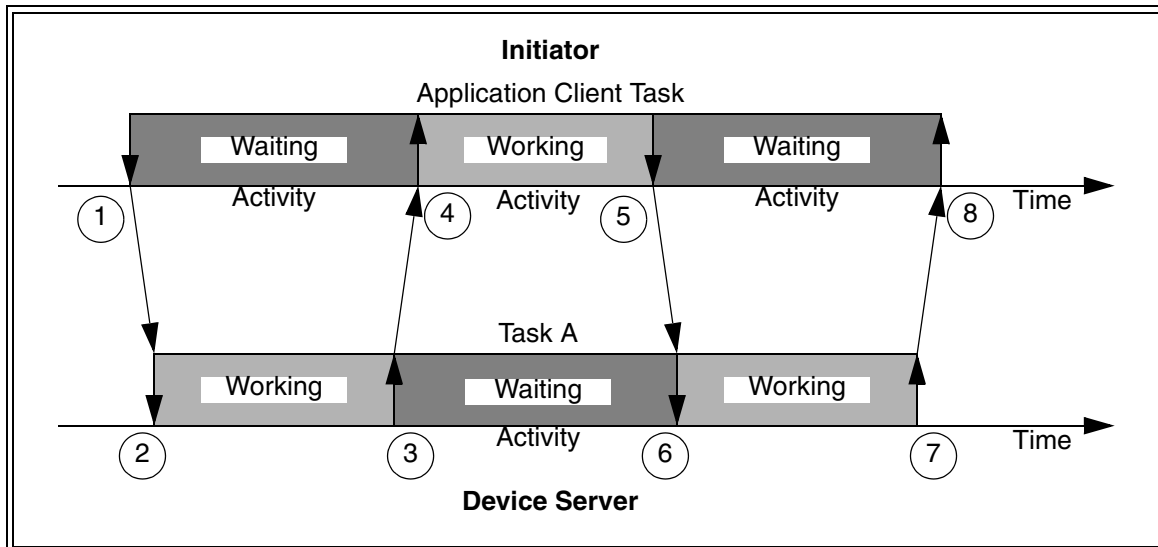


Figure 34 — Linked command processing events

The numbers in figure 34 identify the events described as follows:

1. The application client task performs an **Execute Command** procedure call by invoking the **Send SCSI Command** SCSI transport protocol service to send the CDB and other input parameters to the logical unit. The LINK bit is set to one in the CDB CONTROL byte (see 5.2).
2. The device server is notified through a **SCSI Command Received** indication containing the CDB and command parameters. A task (Task A) is created and entered into the task set.
3. Upon completion of the first command, the device server invokes the **Send Command Complete** SCSI transport protocol service with the status set to INTERMEDIATE or INTERMEDIATE-CONDITION MET and a service response of LINKED COMMAND COMPLETE. Task A is not terminated.
4. The SCSI initiator port returns the status and service response to the application client task by means of a **Command Complete Received** confirmation.
5. The application client task performs an **Execute Command** procedure call by means of the **Send SCSI Command** SCSI transport protocol service as described in step 1. The Task Attribute argument is omitted. The LINK bit in the CDB CONTROL byte is set to zero.
6. The device server receives the last command in the series and processes the operation.
7. The command completes successfully. Task A is terminated. A **Send Command Complete** SCSI transport protocol service response of TASK COMPLETE, with status GOOD, is sent to the application client.
8. The SCSI initiator port delivers an **Command Complete Received** confirmation containing the service response and status to the application client task.

5.9 Command processing considerations and exception conditions

5.9.1 Commands that complete with CHECK CONDITION status

5.9.1.1 Overview

When a command completes with a CHECK CONDITION status, the application client may request that device server alter command processing by establishing an ACA condition. The application client requests establishment of an ACA condition by setting NACA bit to one in the CONTROL byte of the CDB that specified the command. The ACA condition is described in 5.9.2. Requirements that apply when the ACA condition is not in effect are described in 5.9.1.2.

When the NACA bit is set to zero in the CONTROL byte of the CDB for a command that returns in a CHECK CONDITION status, the device server shall not establish an ACA condition. However, tasks other than the task for the command returning the CHECK CONDITION status may be aborted as described in 5.9.1.3.

5.9.1.2 Handling tasks when ACA is not in effect

Table 22 describes the handling of tasks when an ACA condition is not in effect for the task set. The number of SCSI initiator ports in the task set is influenced by the TST field in the Control mode page (see SPC-3).

Table 22 — Task handling when ACA is not in effect

New Task Properties		Device Server Action	ACA Established if New Task Terminates with a CHECK CONDITION status
Attribute ^a	NACA Value ^b		
Any Attribute Except ACA	0	Process the task. ^c	No
	1		Yes
ACA	0	Process an invalid task attribute condition as described in 5.9.5.	No
	1		Yes

^a Task attributes are described in 8.6.
^b The NACA bit is in the CONTROL byte in the CDB (see 5.2).
^c All the conditions that affect the processing of commands (e.g., reservations) still apply.

5.9.1.3 Aborting other tasks when CHECK CONDITION status is returned without establishing an ACA

When a CHECK CONDITION status is returned for a command where the NACA bit is set to zero in the command's CDB CONTROL byte (i.e., when an ACA condition is not established), tasks in the dormant or enabled task state (see 8.5) may be aborted based on the contents of the TST and QERR field in the Control mode page (see SPC-3) as shown in table 23. The TST (task set type) Control mode page field specifies the type of task set in the logical unit (see SPC-3). The QERR (queue error management) Control mode page field specifies how the device server handles blocked and dormant tasks when another task receives a CHECK CONDITION status (see SPC-3).

Table 23 — Aborting tasks when an ACA is not established

QERR	TST	Action
00b	000b	Tasks other than the task returning CHECK CONDITION status shall not be aborted.
	001b	
01b	000b	All enabled and dormant tasks from all SCSI initiator ports shall be aborted (see 5.7).
	001b	All enabled and dormant tasks from the faulted initiator port shall be aborted (see 5.7). All tasks from SCSI initiator ports other than the faulted initiator port shall not be aborted.
11b	000b	All enabled and dormant tasks from the faulted initiator port shall be aborted (see 5.7). All tasks from SCSI initiator ports other than the faulted initiator port shall not be aborted.
	001b	

5.9.2 Auto contingent allegiance (ACA)

When a command is terminated with a CHECK CONDITION status, the application client may request that device server alter command processing by establishing an ACA condition. When the NACA bit in the CONTROL byte of the CDB for the command is set to one, the device server shall establish an ACA condition as described in 5.9.2.1. Upon establishment of the ACA condition, some tasks other than the task returning the CHECK CONDITION status may be aborted and continued processing of other tasks may be blocked as described in 5.9.2.1.

While the ACA condition is in effect, new tasks received by the logical unit from the faulted initiator port are not allowed to enter the task set unless they have the ACA task attribute (see 8.6.5). One of the results of the ACA task attribute requirement is that commands in-flight when the CHECK CONDITION status occurs are returned unprocessed to the SCSI initiator port with an ACA ACTIVE status. Multiple commands may be sent one at a time using the ACA task attribute to recover from the event that resulted in the ACA condition without clearing the ACA.

While the ACA condition is in effect:

- a) New tasks from the faulted initiator port shall be handled as described in 5.9.2.2, and
- b) New tasks from SCSI initiator ports other than the faulted initiator port shall be handled as described in 5.9.2.3.

The methods for clearing an ACA condition are described in 5.9.2.4.

5.9.2.1 Establishing an ACA

When a device server terminates a command with a CHECK CONDITION status and the NACA bit was set to one in the CONTROL byte of the faulting command, the device server shall create an ACA condition.

When an ACA condition is established, tasks in the dormant or enabled task state (see 8.5) shall either be aborted or blocked based on the contents of the TST and QERR field in the Control mode page (see SPC-3) as shown in table 24. The TST (task set type) Control mode page field specifies the type of task set in the logical unit (see

SPC-3). The QE_{ERR} (queue error management) Control mode page field specifies how the device server handles blocked and dormant tasks when another task receives a CHECK CONDITION status (see SPC-3).

Table 24 — Blocking and aborting tasks when an ACA is established

QE _{ERR}	TST	Action
00b	000b	All enabled tasks from all SCSI initiator ports shall transition to the blocked task state (see 8.8). All dormant tasks from all SCSI initiator ports shall remain in the dormant task state.
	001b	All enabled tasks from the faulted initiator port shall transition to the blocked task state (see 8.8). All dormant tasks from the faulted initiator port shall remain in the dormant task state. All tasks from SCSI initiator ports other than the faulted initiator port shall not be affected by the establishment of this ACA condition.
01b	000b	All enabled and dormant tasks from all SCSI initiator ports shall be aborted (see 5.7).
	001b	All enabled and dormant tasks from the faulted initiator port shall be aborted (see 5.7). All tasks from SCSI initiator ports other than the faulted initiator port shall not be affected by the establishment of this ACA condition.
11b	000b	All enabled and dormant tasks from the faulted initiator port shall be aborted (see 5.7). All enabled tasks from SCSI initiator ports other than the faulted initiator port shall transition to the blocked task state (see 8.8). All dormant tasks from SCSI initiator ports other than the faulted initiator port shall remain in the dormant task state.
	001b	All enabled and dormant tasks from the faulted initiator port shall be aborted (see 5.7). All tasks from SCSI initiator ports other than the faulted initiator port shall not be affected by the establishment of this ACA condition.

An ACA condition shall not cross task set boundaries and shall be preserved until it is cleared as described in 5.9.2.4.

If the SCSI transport protocol does not enforce state synchronization as described in 4.6.2, there may be a time delay between the occurrence of the ACA condition and the time at which the application client becomes aware of the condition.

5.9.2.2 Handling new tasks from the faulted initiator port when ACA is in effect

Table 25 describes the handling of new tasks from the faulted initiator port when ACA is in effect.

Table 25 — Handling for new tasks from a faulted initiator port during ACA

New Task Properties		ACA Task Present in the Task Set	Device Server Action	ACA Established If New Task Terminates with a CHECK CONDITION status
Attribute ^a	NACA Value ^b			
ACA	0	No	Process the task. ^d	No ^c
	1	No		Yes ^c
	0 or 1	Yes	Terminate the task with ACA ACTIVE status.	n/a
Any Attribute Except ACA	0 or 1	n/a	Terminate the task with ACA ACTIVE status.	n/a

^a Task attributes are described in 8.6.
^b The NACA bit is in the CONTROL byte in the CDB (see 5.2).
^c If a task with the ACA attribute terminates with a CHECK CONDITION status, the existing ACA condition shall be cleared and the value of the NACA bit shall control the establishment of a new ACA condition.
^d All the conditions that affect the processing of commands (e.g., reservations) still apply.

5.9.2.3 Handling new tasks from non-faulted initiator ports when ACA is in effect

5.9.2.3.1 Commands permitted from non-faulted initiator ports during ACA

The device server shall process a PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action (see SPC-3) while an ACA condition is established when the command is received from a SCSI initiator port other than the faulted initiator port.

NOTE 8 - The processing of specific commands (e.g., PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action) the from SCSI initiator ports other than the faulted initiator port while an ACA condition is in effect provides SCSI initiator ports other than the faulted initiator port the opportunity to recover from error conditions that the faulted initiator port is unable to recover from itself.

5.9.2.3.2 Handling new tasks from non-faulted initiator ports when ACA is in effect

The handling of tasks created by SCSI initiator ports other than the faulted initiator port depends on the value in the TST field in the Control mode page (see SPC-3).

Table 26 describes the handling of new tasks from SCSI initiator ports other than the faulted initiator port when ACA is in effect.

Table 26 — Handling for new tasks from non-faulted initiator ports during ACA

TST Field Value in Control mode page	New Task Properties		New Command Permitted During ACA ^c	Device Server Action	ACA Established If New Task Terminates with a CHECK CONDITION status
	Attribute ^a	NACA Value ^b			
000b	ACA	n/a	n/a	Terminate the task with ACA ACTIVE status.	n/a
	Any Attribute Except ACA	0	No	Terminate the task with BUSY status.	n/a
		1	No	Terminate the task with ACA ACTIVE status.	n/a
		0	Yes	Process the task.	No ^d
		1	Yes		Yes ^d
001b	ACA	0	n/a	Process an invalid task attribute condition as described in 5.9.5.	No
		1		Yes	
	Any Attribute Except ACA	0 or 1	n/a	Process the task. ^e	See 5.9.1.2.

^a Task attributes are described in 8.6.
^b The NACA bit is in the CONTROL byte in the CDB (see 5.2).
^c See 5.9.2.3.1.
^d If a permitted command terminates with a CHECK CONDITION status, the existing ACA condition shall be cleared and the value of the NACA bit shall control the establishment of a new ACA condition.
^e When the TST field in the Control mode page contains 001b, commands from SCSI initiator ports other than the faulted initiator port shall be processed as if the ACA condition does not exist (see 5.9.1.2). In this case, the logical unit shall be capable of handling concurrent ACA conditions and sense data for all SCSI initiator ports.

5.9.2.4 Clearing an ACA condition

An ACA condition shall only be cleared:

- a) As a result of a reset condition (see 6.3.2), logical unit reset (see 6.3.3), or I_T nexus loss (see 6.3.4);
- b) By a CLEAR ACA task management function (see 7.4) from the faulted initiator port;
- c) By a PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action with the ACA task attribute from the faulted initiator port that clears the tasks of the faulted initiator port (see SPC-3);

- d) By a PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action with a task attribute other than ACA from a SCSI initiator port other than the faulted initiator port that clears the tasks of the faulted initiator port;
- e) When a command with the ACA task attribute from the faulted initiator port terminates with a CHECK CONDITION status; or
- f) When a PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action terminates in a CHECK CONDITION status.

Cases e) and f) may result in the establishment of a new ACA based on the value of the NACA bit.

When an ACA condition is cleared and no new ACA condition is established, the state of all tasks in the task set shall be modified as described in 8.8.

5.9.3 Overlapped commands

An overlapped command occurs when a task manager detects the use of a duplicate I_T_L_Q nexus (see 4.11) in a command before a task holding that I_T_L_Q nexus completes its task lifetime (see 5.5). Each SCSI transport protocol standard shall specify whether or not a task manager is required to detect overlapped commands.

A task manager that detects an overlapped command shall abort all tasks for the faulted initiator port in the task set and the device server shall return CHECK CONDITION status for that command. The sense key shall be set to ABORTED COMMAND and the additional sense code shall be set to OVERLAPPED COMMANDS ATTEMPTED.

NOTES

- 9 An overlapped command may be indicative of a serious error and, if not detected, may result in corrupted data. This is considered a catastrophic failure on the part of the SCSI initiator device. Therefore, vendor specific error recovery procedures may be required to guarantee the data integrity on the medium. The SCSI target device logical unit may return additional sense data to aid in this error recovery procedure (e.g., sequential-access devices may return the residue of blocks remaining to be written or read at the time the second command was received).
- 10 Some logical units may not detect an overlapped command until after the CDB has been received.

5.9.4 Incorrect logical unit selection

The SCSI target device's response to an incorrect logical unit number is described in this subclause.

In response to an REQUEST SENSE command or an INQUIRY command the SCSI target device shall respond as defined in SPC-3.

Any command except REQUEST SENSE or INQUIRY:

- a) Shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and with the additional sense code set to LOGICAL UNIT NOT SUPPORTED, if:
 - A) The SCSI target device does not support the logical unit (e.g., some SCSI target devices support only one peripheral device); or
 - B) The SCSI target device supports the logical unit, but the peripheral device is not currently attached to the SCSI target device;
- or
- b) Is responded to in a vendor specific manner, if:
 - A) The SCSI target device supports the logical unit and the peripheral device is attached, but the peripheral device is not operational; or
 - B) The SCSI target device supports the logical unit but is incapable of determining if the peripheral device is attached or is not operational when the peripheral device is not ready.

5.9.5 Task attribute exception conditions

If a command is received with a task attribute that is not supported or is not valid (e.g., an ACA task attribute when an ACA condition does not exist), the command shall be terminated with CHECK CONDITION status, sense key of ILLEGAL REQUEST and additional sense code of INVALID MESSAGE ERROR.

NOTE 11 - The use of the INVALID MESSAGE ERROR additional sense code is based on its similar usage in previous versions of this standard. The use of the INVALID MESSAGE ERROR additional sense code is not to be interpreted as a description of how the task attributes are represented by any given SCSI transport protocol.

Task attribute support should be reported with the Extended INQUIRY Data VPD page (see SPC-3).

5.9.6 Sense data

Sense data shall be made available by the logical unit in the event a command completes with a CHECK CONDITION status or other conditions. The format, content, and conditions under which sense data shall be prepared by the logical unit are specified in this standard, SPC-3, the applicable command set standard and applicable SCSI transport protocol standard.

Sense data shall be preserved by the logical unit for the SCSI initiator port until:

- a) It is transferred;
- b) A logical unit reset (see 6.3.3) occurs; or
- c) An I_T nexus loss (see 6.3.4) involving the SCSI initiator port occurs.

When a command completes with a CHECK CONDITION status, sense data shall be returned in the same I_T_L_Q nexus transaction (see 3.1.46) as the CHECK CONDITION status. After the sense data is returned, it shall be cleared except when it associated with an unit attention condition and the UA_INTLCK_CTRL field in the Control mode page (see SPC-3) contains 10b or 11b.

The return of sense data in the same I_T_L_Q nexus transaction (see 3.1.46) as a CHECK CONDITION status shall not affect ACA (see 5.9.2) or the sense data associated with an unit attention condition when the UA_INTLCK_CTRL field contains 10b or 11b.

5.9.7 Unit Attention condition

Each logical unit shall generate an unit attention condition whenever one of the following events occurs:

- a) A hard reset condition (see 6.3.2), logical unit reset condition (see 6.3.3), or I_T nexus loss condition (see 6.3.4) occurs;
- b) A removable medium may have been changed;
- c) The mode parameters in effect for this SCSI initiator port have been changed by another SCSI initiator port (see SPC-3);
- d) The version or level of microcode has been changed (see SPC-3);
- e) Tasks for this SCSI initiator port were cleared by another SCSI initiator port;
- f) INQUIRY data has been changed (see SPC-3);
- g) The logical unit inventory has been changed (see SPC-3);
- h) The mode parameters in effect for the SCSI initiator port have been restored from non-volatile memory (see SPC-3);
- i) A change in the condition of a synchronized spindle; or
- j) Any other event requiring the attention of the SCSI initiator device.

Logical units may queue unit attention conditions. After the first unit attention condition is cleared, another unit attention condition may exist (e.g., a unit attention condition with an additional sense code of POWER ON OCCURRED followed by one with an additional sense code of MICROCODE HAS BEEN CHANGED).

An unit attention condition shall persist on the logical unit for each SCSI initiator port until that SCSI initiator port clears the condition as described in the remainder of this subclause.

If an INQUIRY command enters the enabled task state, the logical unit shall perform the INQUIRY command and shall neither report nor clear any unit attention condition.

If a REPORT LUNS command enters the enabled task state, the logical unit shall perform the REPORT LUNS command and shall not report any unit attention condition. The logical unit shall clear any unit attention condition established in response to a change in the logical unit inventory for all logical units for the SCSI initiator port that sent the REPORT LUNS command. The logical unit shall not clear any other unit attention condition.

If a REQUEST SENSE command enters the enabled task state while an unit attention condition exists for the SCSI initiator port that sent the REQUEST SENSE command, then the logical unit shall return GOOD status and either:

- a) Report any pending sense data as parameter data and preserve all unit attention conditions on the logical unit; or,
- b) Report an unit attention condition for the SCSI initiator port that sent the REQUEST SENSE command as parameter data. The logical unit may discard any pending sense data and shall clear the reported unit attention condition for that SCSI initiator port.

If the logical unit has already generated the ACA condition for an unit attention condition, the logical unit shall report the unit attention condition (i.e., option b) above).

If a command other than INQUIRY, REPORT LUNS, or REQUEST SENSE enters the enabled task state while an unit attention condition exists for the SCSI initiator port that sent the command, the logical unit shall terminate the command with a CHECK CONDITION status. The logical unit shall provide sense data that reports an unit attention condition for the SCSI initiator port that sent the command.

If a logical unit reports an unit attention condition with a CHECK CONDITION status and the UA_INTLCK_CTRL field in the Control mode page contains 00b (see SPC-3), then the logical unit shall clear the reported unit attention condition for that SCSI initiator port on the logical unit. If the UA_INTLCK_CTRL field in the Control mode page contains 10b or 11b, the logical unit shall not clear unit attention conditions reported with a CHECK CONDITION status.

6 SCSI events and event notification model

6.1 SCSI events overview

SCSI events may occur or be detected in either:

- a) The SCSI device;
- b) One or more SCSI ports within a SCSI device; or
- c) The application client, task manager, or device server.

The detection of any event may require processing by the object that detects it.

Events that occur in the SCSI device are assumed to be detected and processed by all objects within the SCSI device.

When a SCSI port detects an event, it shall use the event notification services (see 6.4) to notify SCSI application layer objects that the event has been detected.

The events detected and event notification services usage depends on whether the SCSI device is a SCSI target device (see figure 35) or a SCSI initiator device (see figure 36). SCSI target/initiator devices shall use the event notification services defined for both SCSI target devices and SCSI initiator devices.

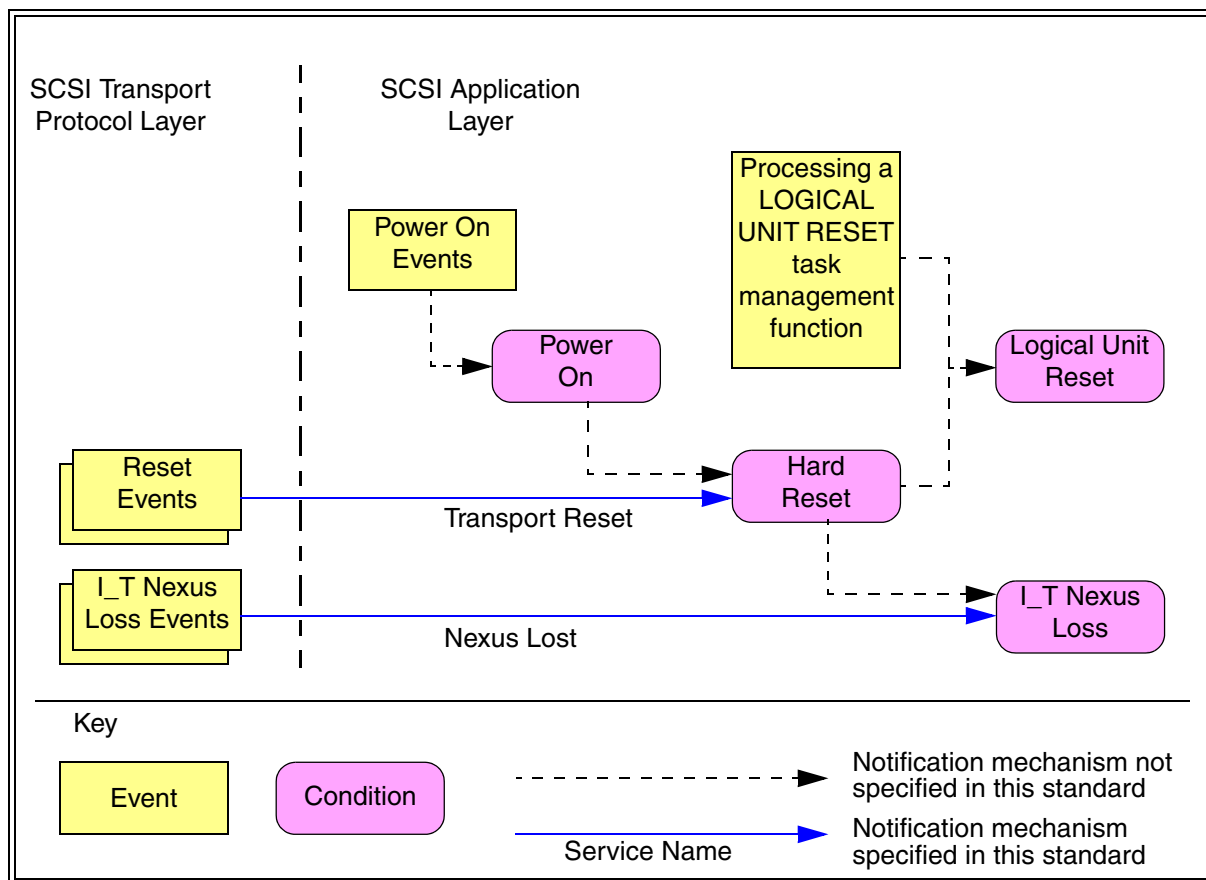


Figure 35 — Events and event notifications for SCSI target devices

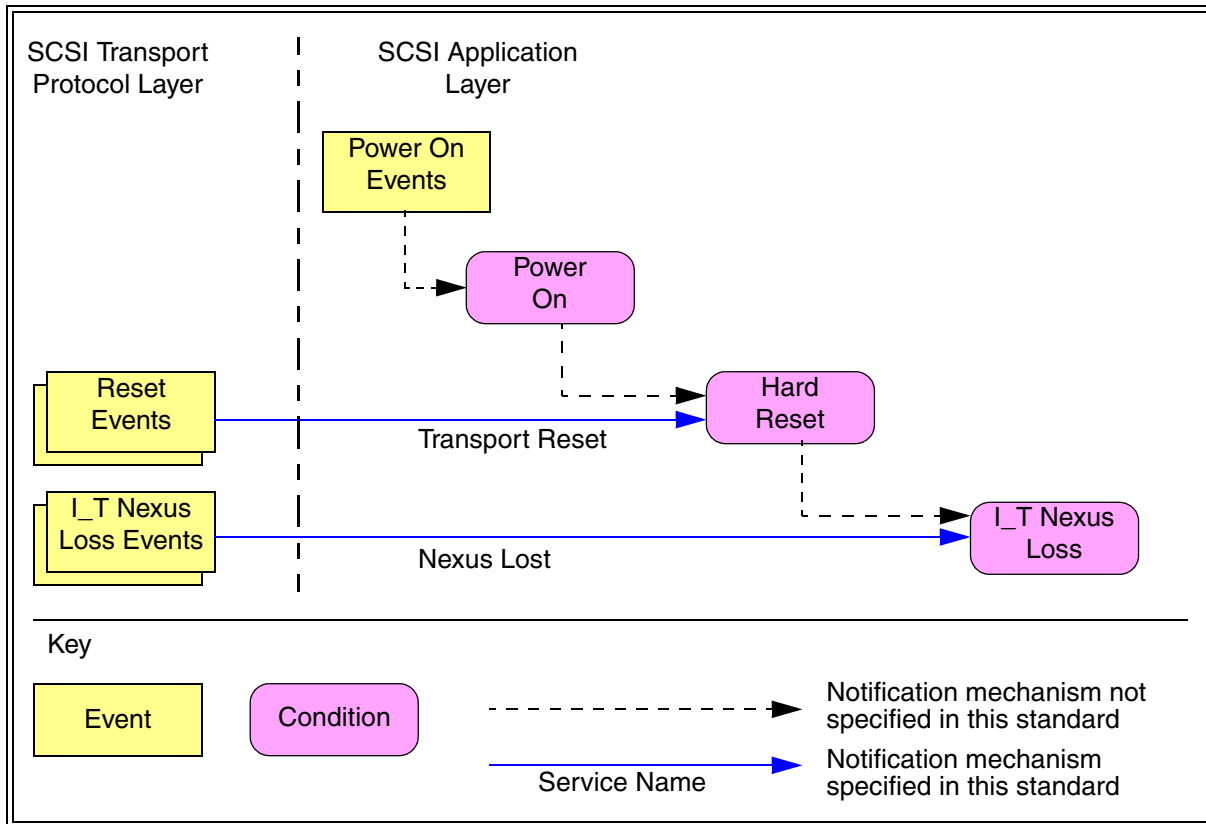


Figure 36 — Events and event notifications for SCSI initiator devices

6.2 Establishing a unit attention condition subsequent to detection of an event

Table 27 shows the additional sense code that a logical unit shall use when a unit attention (see 5.9.7) is established for each of the conditions shown in figure 35 (see 6.1). A SCSI transport protocol may define a more specific additional sense code than SCSI BUS RESET OCCURRED for reset events. The most specific condition in table 27 known to the logical unit should be used to establish the additional sense code for a unit attention.

Table 27 — Unit attention additional sense codes for events detected by SCSI target devices

Condition	Additional Sense Code	Specificity
Logical unit is unable to distinguish between the conditions	POWER ON, RESET, OR BUS DEVICE RESET OCCURRED	Lowest
Power on	POWER ON OCCURRED or DEVICE INTERNAL RESET	
Hard reset	SCSI BUS RESET OCCURRED or protocol specific	
Logical unit reset	BUS DEVICE RESET FUNCTION OCCURRED	
I_T nexus loss	I_T NEXUS LOSS OCCURRED	Highest

NOTE 12 - The names of the unit attention conditions listed in the subclause (e.g., SCSI BUS RESET OCCURRED) are based on usage in previous versions of this standard. The use of these unit attention condition names is not to be interpreted as a description of how the unit attention conditions are represented by any given SCSI transport protocol.

A logical unit may use the I_T NEXUS LOSS OCCURRED additional sense code when establishing a unit attention condition if:

- a) The SCSI initiator port to which the sense data is being delivered is the SCSI initiator port that was associated with the I_T nexus loss, and the logical unit has maintained all state information specific to that SCSI initiator port since the I_T nexus loss; and
- b) The I_T nexus being used to deliver the sense data is the same I_T nexus that was lost, and the logical unit has maintained all state information specific to that I_T nexus since the I_T nexus loss.

Otherwise, the logical unit shall use one of the less specific additional sense codes (e.g., POWER ON OCCURRED) when establishing a unit attention condition.

6.3 Conditions resulting from SCSI events

6.3.1 Power on

Power on is a SCSI device condition resulting from a power on event. When a SCSI device is powered on, it shall cause a hard reset.

The power on condition applies to both SCSI initiator devices and SCSI target devices.

6.3.2 Hard reset

A hard reset is:

- a) A response to a **Transport Reset** event notification (see 6.4); or
- b) A part of the response to a power on condition (see 6.3.1).

The definition of reset events and the notification of their detection is SCSI transport protocol specific.

Each SCSI transport protocol standard that defines reset events shall specify a SCSI target port's protocol specific actions in response to reset events. Each SCSI transport protocol standard that defines reset events should specify when those events result in the delivery of a **Transport Reset** event notification to the SCSI applications layer.

SCSI transport protocols may include reset events that have no SCSI effects (e.g., a Fibre Channel non-initializing loop initialization primitive).

The hard reset condition applies to both SCSI initiator devices and SCSI target devices.

A SCSI target port's response to a hard reset condition shall include a logical unit reset condition (see 6.3.3) for all logical units to which the SCSI target port has access. A hard reset condition shall not affect any other SCSI target ports in the SCSI target device, however, the logical unit reset condition established by a hard reset may affect tasks that are communicating via other SCSI target ports.

Although the task manager response to task management requests is subject to the presence of access restrictions, as managed by ACCESS CONTROL OUT commands (see SPC-3), a hard reset condition shall not be prevented by access controls.

When a SCSI initiator port detects a hard reset condition, it should terminate all its outstanding **Execute Command** procedure calls with a service response of SERVICE DELIVERY OR TARGET FAILURE. A hard reset condition shall not affect any other SCSI initiator ports in the SCSI initiator device, however, the logical unit reset condition

established in a SCSI target device by a hard reset may affect tasks that are communicating via other SCSI initiator ports.

A SCSI port's response to a hard reset condition shall include an I_T nexus loss for every I_T nexus associated with that SCSI port.

6.3.3 Logical unit reset

A logical unit reset is:

- a) The response to a LOGICAL UNIT RESET task management request (see 7.6); or
- b) A part of the response to a hard reset condition (see 6.3.2).

The logical unit reset condition applies only to SCSI target devices.

When responding to a logical unit reset condition, the logical unit shall:

- a) Abort all tasks as described in 5.7;
- b) Clear an ACA condition (see 5.9.2.4), if one is present;
- c) Establish an unit attention condition (see 5.9.7 and 6.2);
- d) Initiate a logical unit reset for all dependent logical units (see 4.14); and
- e) Perform any additional functions required by the applicable command set standards.

6.3.4 I_T nexus loss

An I_T nexus loss event notification is delivered by a SCSI port in response to:

- a) A hard reset; or
- b) an I_T nexus loss event (e.g., logout).

An I_T nexus loss notification event is an indication from the SCSI transport protocol to the SCSI application layer that an I_T nexus no longer exists. SCSI transport protocols may define I_T nexus loss events.

The I_T nexus loss condition applies to both SCSI initiator devices and SCSI target devices.

When a SCSI target port detects an I_T nexus loss, a **Nexus Loss** event notification indication (see 6.4) shall be delivered to each logical unit to which the I_T nexus has access. In response to the resulting I_T nexus loss condition a logical unit shall take the following actions:

- a) Abort all tasks from the SCSI initiator port associated with the I_T nexus as described in 5.7;
- b) Clear an ACA condition (see 5.9.2.4), if one is present for the SCSI initiator port;
- c) Establish an unit attention condition for the SCSI initiator port associated with the I_T nexus (see 5.9.7 and 6.2); and
- d) Perform any additional functions required by the applicable command set standards.

If the logical unit retains state information for the I_T nexus that is lost, its response to the subsequent I_T nexus re-establishment for the logical unit should include establishing a unit attention with an additional sense code of I_T NEXUS LOSS OCCURRED.

If the logical unit does not retain state information for the I_T nexus that is lost, it shall consider the subsequent I_T nexus re-establishment, if any, as the formation of a new I_T nexus for which there is no past history (e.g., establish a unit attention with an additional sense code of POWER ON OCCURRED).

When a SCSI initiator port detects an I_T nexus loss, it should terminate all its outstanding **Execute Command** procedure calls to the SCSI target port associated with the I_T nexus with a service response of SERVICE DELIVERY OR TARGET FAILURE.

6.4 Event notification SCSI transport protocol services

The SCSI transport protocol services described in this subclause are used by a SCSI initiator port or a SCSI target port to deliver an indication to the SCSI application layer that a SCSI event has been detected.

All SCSI transport protocol standards should define the SCSI transport protocol specific requirements for implementing the **Nexus Loss** indication and the **Transport Reset** indication described in this subclause and when these indications are to be delivered to the SCSI applications layer.

The **Nexus Loss** and the **Transport Reset** indications are defined for both SCSI target devices and SCSI initiator devices.

Indication delivered to device servers and application clients:

Nexus Loss (IN (I_T Nexus))

Argument description:

I_T Nexus: The specific I_T nexus that has been detected as lost.

Indication delivered to device servers and application clients:

Transport Reset (IN (SCSI Port))

Argument descriptions:

SCSI Port: The specific SCSI port in the SCSI device for which a transport reset was detected.

7 Task management functions

7.1 Introduction

An application client requests the processing of a task management function by invoking the SCSI transport protocol services described in 7.8, the collective operation of which is modeled in the following procedure call:

Service Response = Function name (IN (nexus))

The task management function names are summarized in table 28.

Table 28 — Task Management Functions

Task Management Function	Nexus	Reference
ABORT TASK	I_T_L_Q	7.2
ABORT TASK SET	I_T_L	7.3
CLEAR ACA	I_T_L	7.4
CLEAR TASK SET	I_T_L	7.5
LOGICAL UNIT RESET	I_T_L	7.6
QUERY TASK	I_T_L_Q	7.7

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12) identifying the task or tasks affected by the task management function.

I_T Nexus: A SCSI initiator port and SCSI target port nexus (see 4.12).

I_T_L Nexus: A SCSI initiator port, SCSI target port, and logical unit nexus (see 4.12).

I_T_L_Q Nexus: A SCSI initiator port, SCSI target port, logical unit, and task tag nexus (see 4.12).

Service Response:

One of the following SCSI transport protocol specific responses shall be returned:

FUNCTION COMPLETE: A task manager response indicating that the requested function is complete. Unless another response is required, the task manager shall return this response upon completion of a task management request supported by the logical unit or SCSI target device to which the request was directed.

FUNCTION SUCCEEDED: An optional task manager response indicating that the requested function is supported and completed successfully. This task manager response shall only be used by functions that require notification of success.

FUNCTION REJECTED: An task manager response indicating that the requested function is not supported by the logical unit or SCSI target device to which the function was directed.

INCORRECT LOGICAL UNIT NUMBER: An optional task manager response indicating that the function requested processing for an incorrect logical unit number.

SERVICE DELIVERY OR TARGET FAILURE: The request was terminated due to a service delivery failure (see 3.1.113) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Each SCSI transport protocol standard shall define the events comprising each of these service responses.

The task manager response to task management requests is subject to the presence of access restrictions, as managed by ACCESS CONTROL OUT and ACCESS CONTROL IN commands (see SPC-3), as follows:

- a) A task management request of ABORT TASK, ABORT TASK SET, CLEAR ACA, or QUERY TASK shall not be affected by the presence of access restrictions;
- b) A task management request of CLEAR TASK SET or LOGICAL UNIT RESET received from a SCSI initiator port that is denied access to the logical unit, either because it has no access rights or because it is in the pending-enrolled state, shall cause no change to the logical unit; and
- c) The task management function service response shall not be affected by the presence of access restrictions.

7.2 ABORT TASK

Request:

Service Response = ABORT TASK (IN (I_T_L_Q Nexus))

Description:

This function shall be supported by all logical units.

The task manager shall abort the specified task, if any, as described in 5.7.2. Previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the ABORT TASK function.

If the logical unit supports this function, a response of FUNCTION COMPLETE shall indicate that the task was aborted or was not in the task set. In either case, the SCSI target device shall guarantee that no further responses from the task are sent to the SCSI initiator port.

All SCSI transport protocol standards shall support the ABORT TASK task management function.

7.3 ABORT TASK SET

Request:

Service Response = ABORT TASK SET (IN (I_T_L Nexus))

Description:

This function shall be supported by all logical units.

The task manager shall abort all tasks in the task set that were created by the SCSI initiator port routed through the SCSI target port as described in 5.7.

The task manager shall perform an action equivalent to receiving a series of ABORT TASK requests. All tasks from that SCSI initiator port in the task set shall be aborted. Tasks from other SCSI initiator ports or in other task sets shall not be aborted. Other previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the ABORT TASK SET function.

All SCSI transport protocol standards shall support the ABORT TASK SET task management function.

7.4 CLEAR ACA

Request

Service Response = CLEAR ACA (IN (I_T_L Nexus))

Description:

This function shall be supported by a logical unit if it supports ACA (see 5.2).

For the CLEAR ACA task management function, the task set shall be the one defined by the TST field in the Control mode page (see SPC-3).

The application client issues CLEAR ACA to clear an ACA condition from the task set serviced by the logical unit. The state of all tasks in the task set shall be modified as described in 8.8. For a task with the ACA attribute (see 8.6.5) receipt of a CLEAR ACA function shall have the same effect as receipt of an ABORT TASK function (see 7.2) specifying that task. If successful, this function shall be terminated with a service response of FUNCTION COMPLETE.

If the task manager clears the ACA condition, any task within that task set may be completed subject to the requirements for task set management specified in clause 8.

All SCSI transport protocol standards shall support the CLEAR ACA task management function.

7.5 CLEAR TASK SET

Request:

Service Response = CLEAR TASK SET (IN (I_T_L Nexus))

Description:

This function shall be supported by logical units supporting the full task management model (see 8.3.2) and may be supported by logical unit supporting the basic task management model (see 8.3.3).

For the CLEAR TASK SET task management function, the task set shall be the one defined by the TST field in the Control mode page (see SPC-3).

All tasks in the task set shall be aborted as described in 5.7. The medium may have been altered by partially processed commands.

All pending status and sense data for the task set shall be cleared. Other previously established conditions, including MODE SELECT parameters, reservations, and ACA shall not be changed by the CLEAR TASK SET function.

All SCSI transport protocol standards shall support the CLEAR TASK SET task management function.

7.6 LOGICAL UNIT RESET

Request:

Service Response = LOGICAL UNIT RESET (IN (I_T_L Nexus))

Description:

This function shall be supported by all logical units.

Before returning a FUNCTION COMPLETE response, the logical unit shall perform the logical unit reset functions specified in 6.3.3.

NOTE 13 - Previous versions of this standard only required LOGICAL UNIT RESET support in logical units that supported hierarchical logical units.

All SCSI transport protocol standards shall support the LOGICAL UNIT RESET task management function.

7.7 QUERY TASK

Request:

Service Response = QUERY TASK (IN (I_T_L_Q Nexus))

Description:

SCSI transport protocols may or may not support QUERY TASK and may or may not require logical units accessible through SCSI target ports using such transport protocols to support QUERY TASK. The task manager shall return a response of FUNCTION SUCCEEDED if the specified task exists, or FUNCTION COMPLETE if the specified task does not exist.

7.8 Task management SCSI transport protocol services

The SCSI transport protocol services described in this subclause are used by a SCSI initiator device and SCSI target device to process a task management procedure call. The following arguments are passed:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

Function Identifier: Argument encoding the task management function to be performed.

All SCSI transport protocol standards shall define the SCSI transport protocol specific requirements for implementing the **Send Task Management Request** SCSI transport protocol service and the **Received Task Management Function Executed** confirmation described below. Support for the **Task Management Request Received** indication and **Task Management Function Executed** SCSI transport protocol service response by the SCSI transport protocol standard is optional. All SCSI devices shall implement these SCSI transport protocol services as defined in the applicable SCSI transport protocol standards.

Request sent by an application client:

Send Task Management Request (IN (Nexus, Function Identifier))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

Function Identifier: Argument encoding the task management function to be performed.

Indication received by the task manager:

Task Management Request Received (IN (Nexus, Function Identifier))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

Function Identifier: Argument encoding the task management function to be performed.

Response from task manager:

Task Management Function Executed (IN (Nexus, Service Response))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

Service Response: An encoded value representing one of the following:

FUNCTION COMPLETE: The requested function has been completed.

FUNCTION SUCCEEDED: The requested function is supported and completed successfully.

FUNCTION REJECTED: The task manager does not implement the requested function.

INCORRECT LOGICAL UNIT NUMBER: An optional task manager response indicating that the function requested processing for an incorrect logical unit number.

SERVICE DELIVERY OR TARGET FAILURE: The request was terminated due to a service delivery failure (see 3.1.113) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Confirmation received by application client:

Received Task Management Function Executed (IN (Nexus, Service Response))

Argument descriptions:

Nexus: An I_T Nexus, I_T_L Nexus, or I_T_L_Q Nexus (see 4.12).

Service Response: An encoded value representing one of the following:

FUNCTION COMPLETE: The requested function has been completed.

FUNCTION SUCCEEDED: The requested function is supported and completed successfully.

FUNCTION REJECTED: The task manager does not implement the requested function.

INCORRECT LOGICAL UNIT NUMBER: An optional task manager response indicating that the function requested processing for an incorrect logical unit number.

SERVICE DELIVERY OR TARGET FAILURE: The request was terminated due to a service delivery failure (see 3.1.113) or SCSI target device malfunction. The task manager may or may not have successfully performed the specified function.

Since the nexus used by all task management functions except ABORT TASK and QUERY TASK does not contain a task tag to uniquely identify the task, there may be no way for an application client to associate a confirmation with a request. A SCSI transport protocol that does not provide such an association should not allow a SCSI initiator port to have more than one pending task management request per I_T_L nexus.

7.9 Task management function example

Figure 37 shows the sequence of events associated with a task management function.




Figure 37 — Task management processing events

The numbers in figure 37 identify the events described as follows:

1. The application client task issues a task management request by invoking the **Send Task Management Request** SCSI transport protocol service.
2. The task manager is notified through a **Task Management Request Received** and begins processing the function.
3. The task manager performs the requested function and responds by invoking the **Task Management Function Executed** SCSI transport protocol service to notify the application client. The service response argument is set to a value of FUNCTION COMPLETE.
4. A **Received Task Management Function Executed** confirmation is received by the application client task.

8 Task set management

8.1 Introduction to task set management

This clause describes some of the controls application clients have over task set management behaviors (see 8.3). This clause also specifies task set management requirements in terms of:

- a) Task states (see 8.5);
- b) Task attributes (see 8.6);
- c) Task priority (see 8.7);
- d) The events that cause transitions between task states (see 8.4 and 8.5); and
- e) A map of task state transitions (see 8.8).

This clause concludes with several task set management examples (see 8.9).

Task behavior, as specified in this clause, refers to the functioning of a task as observed by an application client, including the results of command processing and interactions with other tasks.

The requirements for task set management only apply to a task after it has been entered into a task set. A task shall be entered into a task set unless:

- a) A condition exists that causes that task to be completed with a status of BUSY, RESERVATION CONFLICT, TASK SET FULL, or ACA ACTIVE;
- b) Detection of an overlapped command (see 5.9.3) causes that task to be completed with a CHECK CONDITION status; or
- c) SCSI transport protocol specific errors cause that task to be completed with a status other than GOOD.

8.2 Implicit head of queue

A command standard (see 3.1.16) may define tasks that may be processed by the task manager as if the task's task attribute is HEAD OF QUEUE without regard to the actual task attribute received with the task.

8.3 Task management models

8.3.1 Task management model management features

The standard INQUIRY data CMDQUE bit and BQUE bit (see SPC-3) indicate support for the full task management model or basic task management model described in this subclause.

The Control mode page (see SPC-3) contains fields that specify particular task set management behaviors described in this clause.

8.3.2 Full task management model

The full task management model requires the following task set management behaviors:

- a) The SIMPLE task attribute (see 8.6.1) shall be supported;
- b) Task attributes other than SIMPLE may be supported;
- c) The QUEUE ALGORITHM MODIFIER field in the Control mode page (see SPC-3) shall control the processing sequence of tasks having the SIMPLE task attribute;

- d) The QERR field in the Control mode page (see SPC-3) shall control aborting of tasks when a CHECK CONDITION status is returned for any task; and
- e) The CLEAR TASK SET task management function (see 7.5) shall be supported.

8.3.3 Basic task management model

The basic task management model requires the following task set management behaviors:

- a) The SCSI transport protocol and task manager support only:
 - A) The SIMPLE task attribute (see 8.6.1); or
 - B) The ORDERED task attribute (see 8.6.3);
- b) If the SIMPLE task attribute is supported, the task manager may reorder the actual processing sequence of tasks in any manner. Any data integrity exposures related to task sequence order shall be explicitly handled by the application client;
- c) All tasks in the task set shall be aborted when a CHECK CONDITION status is returned for any task (i.e., they shall be handled as if the QERR field in the Control mode page (see SPC-3) is set to 01b);
- d) If the Control mode page is supported, the QUEUE ALGORITHM MODIFIER field shall be set to 1h and the QERR field shall be set to 01b and they shall not be changeable; and
- e) The CLEAR TASK SET task management function (see 7.5) may be supported.

8.4 Task management events

The following describe the events that cause changes in task state.

All older tasks ended: If the TST field in the Control mode page (see SPC-3) equals 000b, all tasks have ended that were accepted from all SCSI initiator ports earlier in time than the referenced task. If the TST field in the Control mode page equals 001b, all tasks have ended that were accepted from the referenced SCSI initiator port earlier in time than the referenced task.

All older head of queue and older ordered tasks ended: If the TST field in the Control mode page equals 000b, all head of queue and ordered tasks have ended that were accepted from all SCSI initiator ports earlier in time than the referenced task. If the TST field in the Control mode page equals 001b, all head of queue and ordered tasks have ended that were accepted from the referenced SCSI initiator port earlier in time than the referenced task.

ACA establishment: An ACA condition has been established (see 5.9.1).

task abort: A task has been aborted as described in 5.7.

task completion: The device server has sent a service response of TASK COMPLETE for the task (see 5.1 and 5.5).

task ended: A task has completed or aborted.

ACA cleared: An ACA condition has been cleared (see 5.9.2.4).

8.5 Task states

8.5.1 Overview

8.5.1.1 Task state nomenclature

This standard defines four task states, summarized in table 29.

Table 29 — Task State Nomenclature

Task State Name	Reference	Tasks in This State May Be Called
Enabled task state	8.5.2	Enabled tasks
Blocked task state	8.5.3	Blocked tasks
Dormant task state	8.5.4	Dormant tasks
Ended task state	8.5.5	Ended tasks

8.5.1.2 Suspended information

Any information the logical unit has or accepts for a task in the blocked task state (see 8.5.3) or dormant task state (see 8.5.4) is required to be held in a condition where it is not available to the task. Such information is called suspended information.

8.5.2 Enabled task state

A task in the enabled task state may become a current task and may complete at any time, subject to the task completion constraints specified in the Control mode page (see SPC-3). A task that has been accepted into the task set shall not complete or become a current task unless it is in the enabled task state.

Except for the use of resources required to preserve task state, a task shall produce no effects detectable by the application client before the task's first transition to the enabled task state. Although, before entering this state for the first time, the task may perform other activities visible at the STPL (e.g., pre-fetching data to be written to the media). This activity shall not result in a detectable change in state as perceived by an application client. In addition, the behavior of a completed task, as defined by the commands it has processed, shall not be affected by the task's states before it enters the enabled task state.

8.5.3 Blocked task state

A task in the blocked task state is prevented from completing due to an ACA condition. A task in this state shall not become a current task. While a task is in the blocked task state, any information the logical unit has or accepts for the task shall be suspended. If the TST field in the Control mode page (see SPC-3) equals 000b the blocked task state is independent of the SCSI initiator port. If the TST field equals 001b the blocked task state applies only to the faulted initiator port.

8.5.4 Dormant task state

A task in the dormant task state is prevented from completing due to the presence of certain other tasks in the task set. A task in this state shall not become a current task. While a task is in the dormant task state, any information the logical unit has or accepts for the task shall be suspended.

8.5.5 Ended task state

A task in the ended task state is removed from the task set.

8.5.6 Task states and task lifetimes

Figure 38 shows the events corresponding to two task processing sequences. Except for the dormant task state between times A and B in case 1, logical unit conditions and the commands processed by the task are identical. Assuming in each case the task completes with a status of GOOD at time C, the state observed by the application client for case 1 shall be indistinguishable from the state observed for case 2.

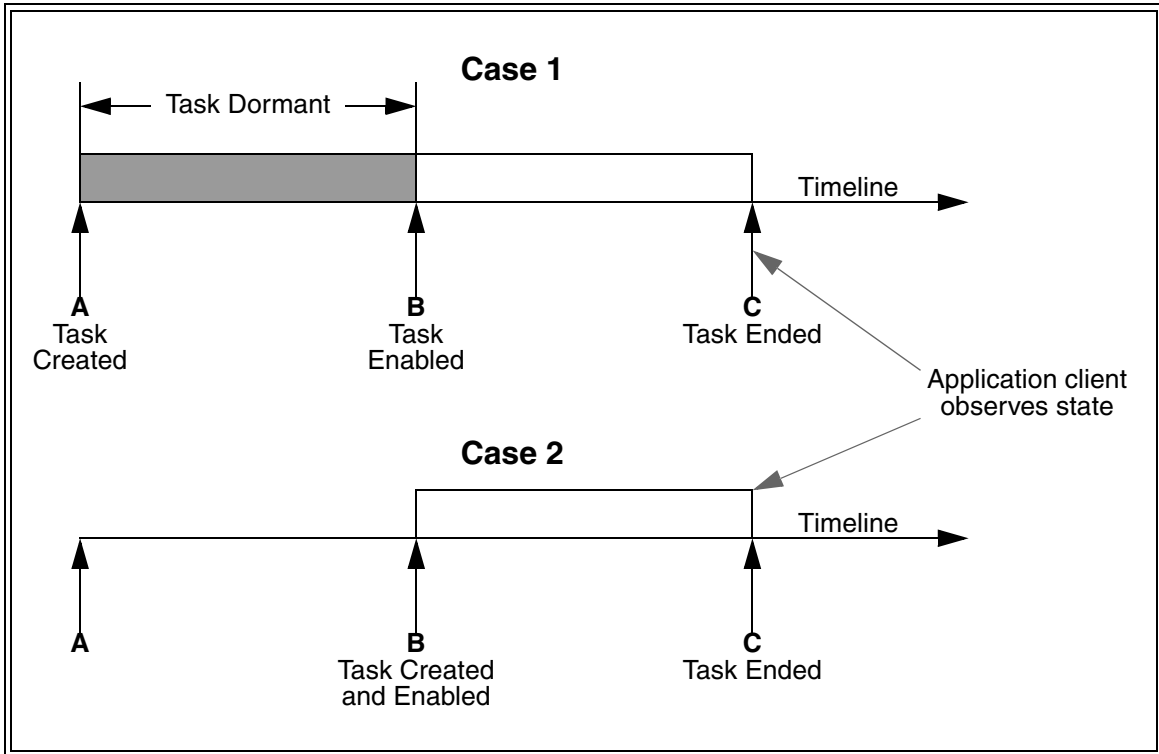


Figure 38 — Example of Dormant state task behavior

8.6 Task attributes

8.6.1 Overview

The application client shall assign a task attribute (see table 30) to each task.

Table 30 — Task attributes

Task Attribute	Reference
SIMPLE	8.6.2
ORDERED	8.6.3
HEAD OF QUEUE	8.6.4
ACA	8.6.5

SCSI transport protocols shall either:

- a) Provide the capability to specify a unique task attribute for each task; or
- b) Ignore the application client assigned task attribute and require use of the basic task management model (see 8.3.3).

SCSI transport protocols should provide the capability to specify a unique task attribute for each task.

8.6.2 Simple task

If accepted, a task having the SIMPLE attribute shall be entered into the task set in the dormant task state. The task shall not enter the enabled task state until all older head of queue tasks and older ordered tasks in the task set have ended (see 8.4).

The QUEUE ALGORITHM MODIFIER field in the Control mode page (see SPC-3) provides additional constraints on task completion order for tasks having the SIMPLE task attribute;

8.6.3 Ordered task

If accepted, a task having the ORDERED attribute shall be entered into the task set in the dormant task state. The task shall not enter the enabled task state until all older tasks in the task set have ended (see 8.4).

8.6.4 Head of queue task

If accepted, a task having the HEAD OF QUEUE attribute shall be entered into the task set in the enabled task state.

8.6.5 ACA task

If accepted, a task having the ACA attribute shall be entered into the task set in the enabled task state. There shall be no more than one ACA task per task set (see 5.9.2.1).

8.7 Task priority

A task priority argument that is set to a value other than zero specifies the relative scheduling importance of a task having a SIMPLE task attribute in relation to other tasks already in the task set. Task priority 1h has the highest scheduling importance, with increasing task priority values indicating lower scheduling importance.

If the task priority argument is set to zero or is not contained within the SCSI transport protocol service indication, a priority assigned to the I_T_L nexus may be used as the task priority. A priority may be assigned to an I_T_L nexus by a SET PRIORITY command (see SPC-3) or by the INITIAL PRIORITY field in the Control Extension mode page (see SPC-3). If no priority has been assigned to the I_T_L nexus using the SET PRIORITY command and the logical unit does not support the INITIAL PRIORITY field of the Control Extension mode page the priority assigned to the task is vendor specific.

A task manager may use task priority to determine an ordering to process tasks with the SIMPLE task attribute within the task set. A difference in task priority between tasks does not necessarily override other scheduling considerations (e.g., different times to access different logical block addresses) or vendor specific scheduling considerations. However, processing of a collection of tasks with different task priorities should cause the subset of tasks with the higher priority to return status sooner in aggregate than the same subset would if the same collection of tasks were submitted under the same conditions but with all task priorities being equal.

For a task that processes linked commands, the task priority shall be that specified for the first command in the series of linked commands. The task priority specified for the second and subsequent commands shall be ignored.

The size of the task priority argument shall be four bits.

8.8 Task state transitions

This subclause describes task state transitions, actions and associated triggering events as they appear to an application client. The logical unit response to events affecting multiple tasks (e.g., a CLEAR TASK SET) may be different from the response to an event affecting a single task. To the application client, the collective behavior appears as a series of state changes occurring to individual tasks.

The task state diagram of figure 39 shows the behavior of a single task in response to an external event.

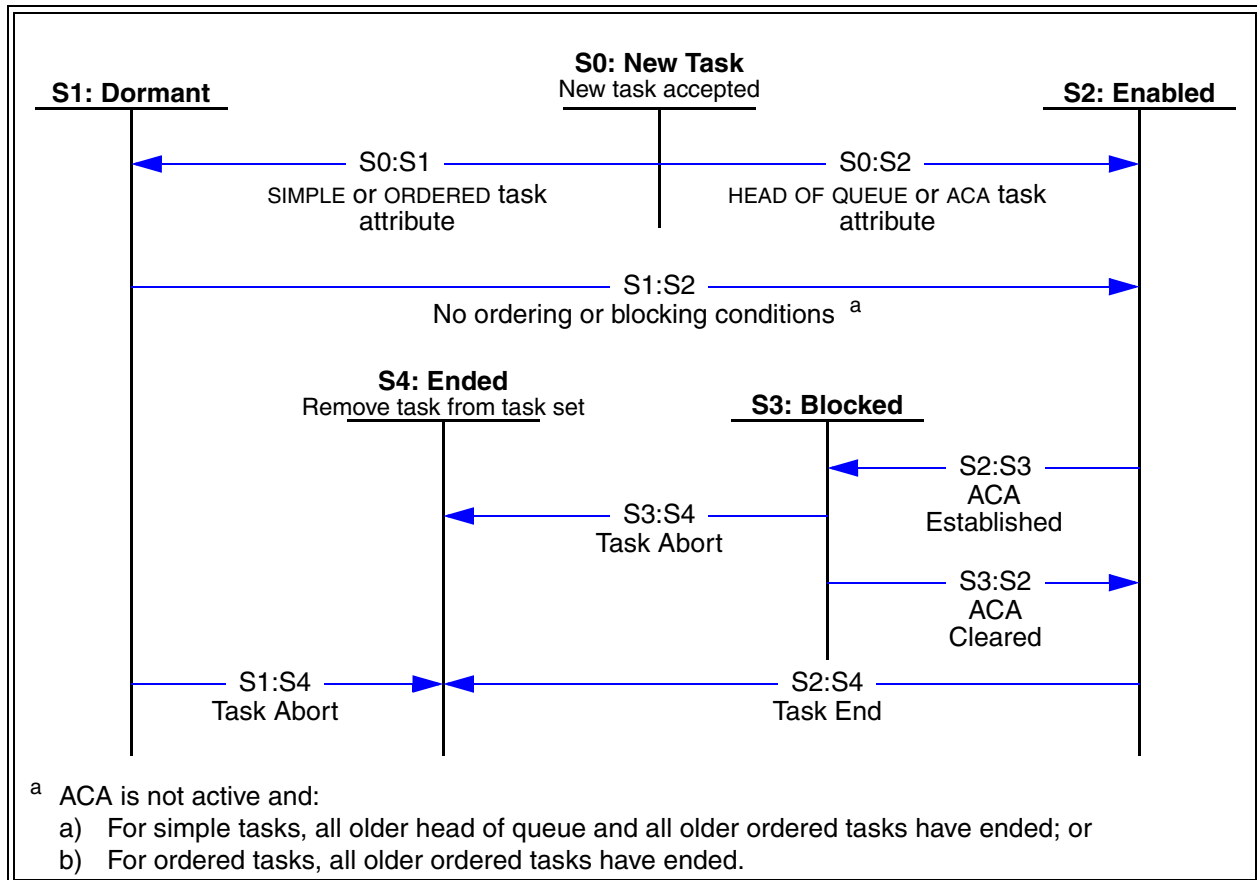


Figure 39 — Task states

Transition S0:S1: If a newly accepted task has the SIMPLE or ORDERED task attribute, it shall transition to the dormant task state.

Transition S0:S2: If a newly accepted task has the HEAD OF QUEUE or ACA task attribute, it shall transition to the enabled task state.

Transition S1:S2: The task attribute of a dormant task shall affect the transition to the enabled task state as follows:

- a) A dormant task having the SIMPLE task attribute shall enter the enabled task state when all older head of queue and older ordered tasks (see 8.4) have ended; or
- b) A dormant task having the ORDERED task attribute shall enter the enabled task state when all older tasks (see 8.4) have ended.

If the TST field in the Control mode page (see SPC-3) contains 000b, then the transition from dormant task to enabled task shall not occur while an ACA is in effect for any SCSI initiator port (see 5.9.2.2 and 5.9.2.3). If the TST field in the Control mode page contains 001b, then dormant tasks from the faulted initiator port shall not transition to the enabled task state while an ACA is in effect for that SCSI initiator port (see 5.9.2.2).

Transition S2:S3: The establishment of an ACA condition (see 8.4) shall cause zero or more enabled tasks to enter the blocked task state as described in 5.9.2.1.

Transition S3:S2: When an ACA condition is cleared (see 8.4), tasks that entered the blocked state when the ACA condition was established (see 5.9.2.1) shall re-enter the enabled task state.

Transition S2:S4: A task that has completed (see 8.4) or aborted (see 8.4 and 5.7) shall enter the ended task state. This is the only state transition that applies to an ACA task.

Transitions S1:S4, S3:S4: A task abort event (see 8.4 and 5.7) shall cause the task to unconditionally enter the ended task state.

8.9 Task set management examples

8.9.1 Introduction

Several task set management scenarios are shown in 8.9.2, 8.9.3, and 8.9.4. The examples are valid for configurations with one or multiple SCSI initiator ports when the TST field contains 000b (i.e., the interaction among tasks in a task set is independent of the SCSI initiator port originating a task). The examples are also valid for a single SCSI initiator port when the TST field contains 001b (i.e., task set management proceeds independently for each SCSI initiator port and the events and transitions in one SCSI initiator port's task set do not affect the task set management for another SCSI initiator port's task set). Throughout these examples, the scope of the task set box drawn in each snapshot depends on the setting of the TST field in the Control mode page (see SPC-3).

The figure accompanying each example shows successive snapshots of a task set after various events (e.g., task creation or completion). In all cases, the constraints on task completion order established using Control mode page (see SPC-3) fields other than the TST field (e.g., the QUEUE ALGORITHM MODIFIER field) are not in effect.

A task set is shown as an ordered list or queue of tasks with the head of the queue towards the top of the figure. A new head of queue task always enters the task set at the head, displacing older head of queue tasks. Simple, ordered and ACA tasks always enter the task set at the end of the queue.

Tasks, denoted by rectangles, are numbered in ascending order from oldest to most recent. Fill, shape and line weight are used to distinguish task states and attributes are shown in table 31.

Table 31 — Task attribute and state indications in examples

Task Attribute	Box Shape	Line Weight	Task State	Box Fill
SIMPLE	Rounded Corners	Thin	Enabled	White
ORDERED	Square Corners	Thin	Dormant	Grey
HEAD OF QUEUE	Square Corners	Thick	Blocked	Black
ACA	Square Corners	Thin Dashed		

The conditions preventing a dormant task from entering enabled task state, except for ACA conditions, are shown by means of blocking boundaries. Such boundaries appear as horizontal lines with an arrow on both ends. The tasks causing the barrier condition are described as part of each example. A task is impeded by the barrier if it is between the boundary and the end of the queue. When ACA is not in effect, a task enters the enabled task state after all intervening barriers have been removed.

8.9.2 Head of queue tasks

Figure 40 shows task set conditions when several head of queue tasks are processed.

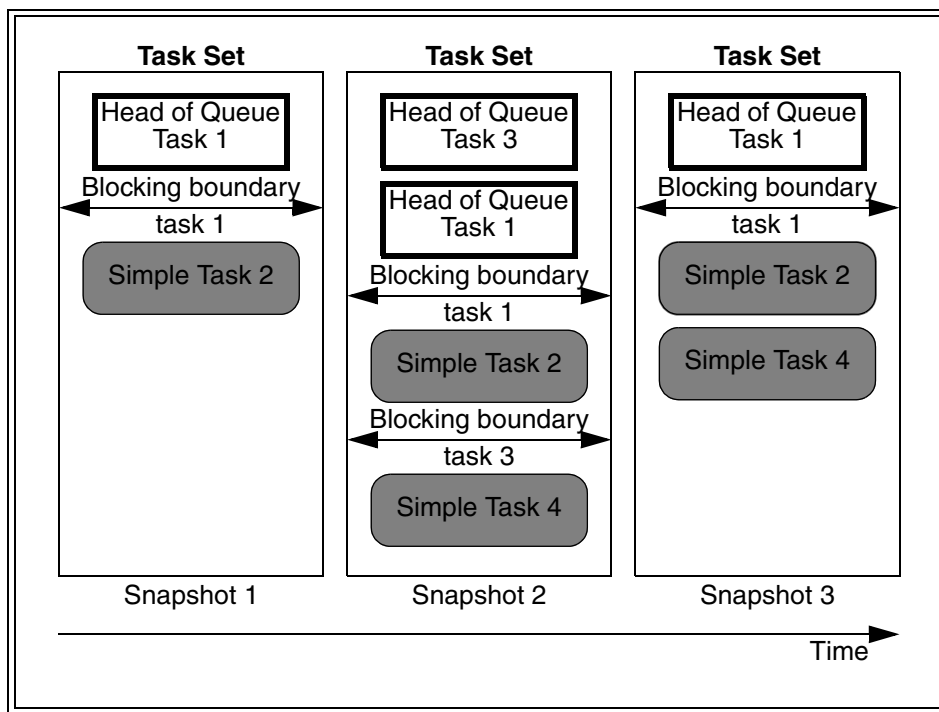


Figure 40 — Head of queue tasks and blocking boundaries (example 1)

In snapshot 1 the task set initially contains one head of queue and one simple task. As shown by the blocking boundary, simple task 2 is in the dormant task state because of the older head of queue task. Snapshot 2 shows the task set after head of queue task 3 and simple task 4 are created. The new head of queue task is placed at the front of the queue in the enabled task state, displacing task 1. Snapshot 3 shows the task set after task 3 completes. Since the conditions indicated by the task 1 blocking boundary are still in effect, tasks 2 and 4 remain in the dormant task state.

Figure 41 is the same as the previous example, except that task 1 completes instead of task 3.

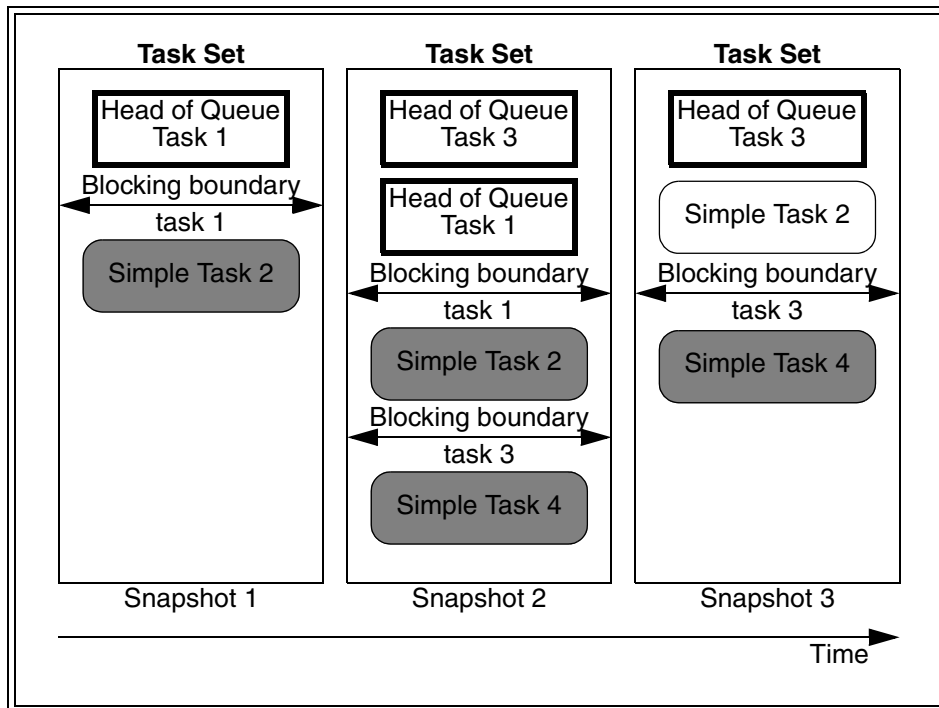


Figure 41 — Head of queue tasks and blocking boundaries (example 2)

The completion of task 1 allows task 2 to enter the enabled task state. Simple task 4 is placed in the dormant task state until task 3 completes.

8.9.3 Ordered tasks

An example of ordered and simple task interaction is shown in figure 42.




Figure 42 — Ordered tasks and blocking boundaries

The state of dormant tasks 2 through 5 is determined by the requirements shown in table 32.

Table 32 — Dormant task blocking boundary requirements

Task	Reason for blocking boundary
2	An ordered task is not allowed to enter the enabled task state until all older tasks have ended.
5	
3	A simple task is not allowed to enter the enabled task state until all older head of queue and older ordered tasks have ended.
4	

The table 32 constraints are shown by the blocking boundaries in snapshot 1.

In snapshot 2, the completion of task 1 allows ordered task 2 to enter the enabled task state. Since the initial constraints on tasks 3, 4 and 5 are still in effect, these tasks are required to remain in the dormant task state. As shown in snapshot 3, the completion of task 2 triggers two state changes: the transitions of task 3 and task 4 to the enabled task state. Task 5 are required to remain in the dormant task state until these tasks end.

8.9.4 ACA task

Figure 43 shows the effects of an ACA condition on the task set. This example assumes the QERR field contains 00b in the Control mode page (see SPC-3). Consequently, clearing an ACA condition does not cause tasks to be aborted.

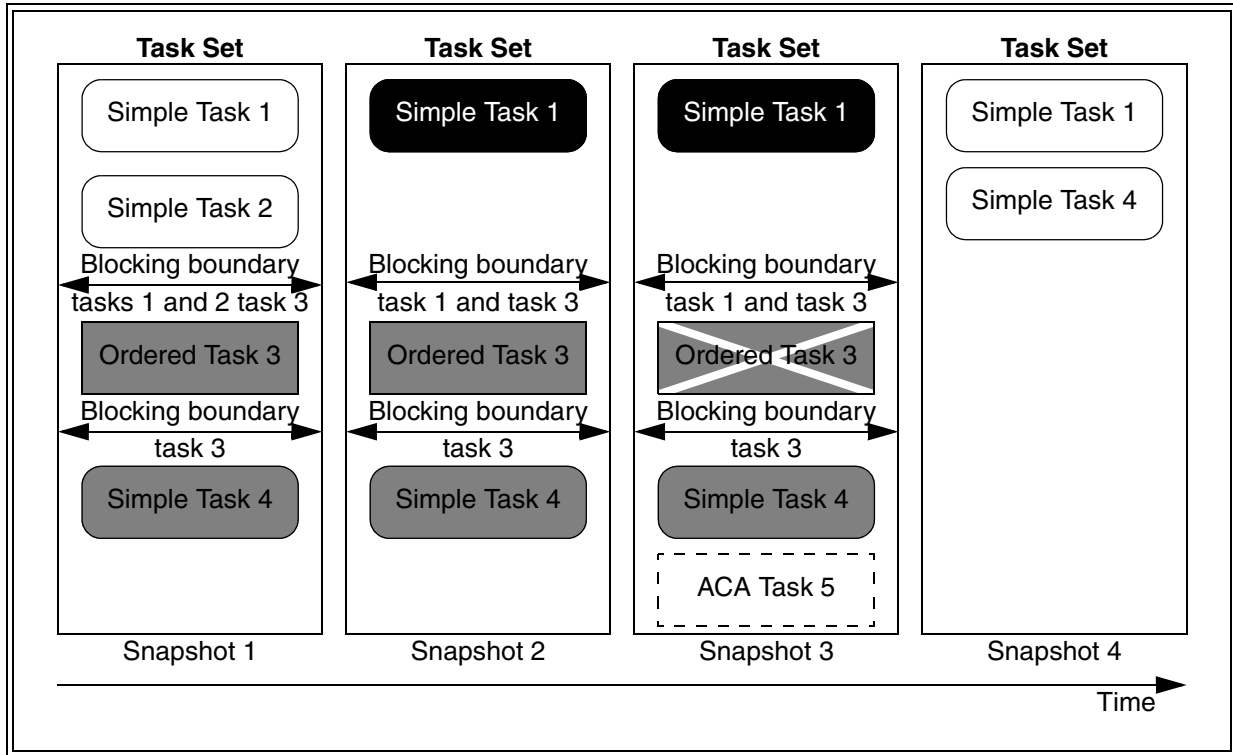


Figure 43 — ACA task example

The completion of task 2 with CHECK CONDITION status causes task 1 to enter the blocked task state shown in snapshot 2. In snapshot 3, ordered task 3 is aborted using the ABORT TASK task management function and ACA task 5 is created to perform additional handling for the exception. Once the ACA condition is cleared (snapshot 4), simple task 1 is allowed to reenter the enabled task state. Since there are no head of queue or ordered tasks older than task 4, it too is allowed enter the enabled task state.

Annex A (informative)

Identifiers and names for objects

A.1 Identifiers and names overview

This annex summarizes SCSI identifiers and names.

The following SCSI architecture model objects have identifiers and names summarized in this annex:

- a) SCSI initiator port (see 3.1.95);
- b) SCSI target port (see 3.1.100);
- c) Logical unit (see 3.1.60);
- d) SCSI initiator device (see 3.1.94); and
- e) SCSI target device (see 3.1.99).

A.2 Identifiers and names

This standard defines the identifiers and names for the objects listed in A.1. The size requirements placed on identifiers by this standard are as shown in table A.1. This standard places no requirements on the sizes of names. Table A.1 also lists whether this standard or SPC-3 requires SCSI transport protocols and logical units to support identifiers and names for an object.

Table A.1 — Object size and support requirements

Object	Identifier		Name	
	Size	Support Requirements	Size	Support Requirements
Initiator device	n/a	n/a	not specified	optional
Target device	n/a	n/a	not specified ^a	optional
Initiator port	not specified	mandatory	not specified	optional
Target port	not specified	mandatory	not specified ^a	optional
Logical unit	8 bytes (maximum)	mandatory	not specified ^a	mandatory

^a Reported in the Device Identification VPD page (see SPC-3).

Each SCSI transport protocol defines the size and format of identifiers and names for each object.

See table A.2 for a list of the size of the identifiers for each SCSI transport protocol.

Table A.2 — Object identifier size for each SCSI transport protocol

Object	Identifier size					
	SPI-5	FCP-2	SRP	iSCSI	SBP-3	SAS SSP
Initiator port	4 bits ^a	3 bytes	16 bytes	246 bytes ^b	2 bytes	8 bytes
Target port	4 bits ^a	3 bytes	16 bytes	233 bytes ^b	11 bytes	8 bytes
Logical unit	6 bits (data group transfers) 8 bytes (information unit transfers)	8 bytes	8 bytes	8 bytes	2 bytes	8 bytes
^a SPI-5 uses a bit significant representation of the SCSI port identifier, therefore, the maximum number of SCSI ports is 16, a value that can be represented in 4 bits. ^b Including the terminating null character byte.						

See table A.3 for a list of the format of the identifiers for each SCSI transport protocol.

Table A.3 — Object identifier format for each SCSI transport protocol

Object	Identifier format					
	SPI-5	FCP-2	SRP	iSCSI	SBP-3	SAS SSP
Initiator port	bit significant (a maximum of 16 ports; one for each bit)	Fibre Channel address identifier	EUI-64 8 byte extension ^a	iSCSI name ^b ",i,0x" Initiator Session Identifier ^c	binary value	NAA IEEE Registered format
Target port	bit significant (a maximum of 16 ports; one for each bit)	Fibre Channel address identifier	EUI-64 8 byte extension ^a	iSCSI name ^b ",i,0x" Target Portal Group Tag ^d	EUI-64 Discovery ID ^e	NAA IEEE Registered format
Logical unit	binary value (6 bit) or as specified in this standard (see 4.9)	as specified in this standard (see 4.9)	as specified in this standard (see 4.9)	as specified in this standard (see 4.9)	binary value (2 bytes)	as specified in this standard (see 4.9)
<p>Key: means "concatenated with"</p> <p>","i,0x" means a UTF-8 string composed of the following five characters: comma, lowercase i, comma, zero, and lowercase x.</p> <p>","t,0x" means a UTF-8 string composed of the following five characters: comma, lowercase t, comma, zero, and lowercase x.</p>						
<p>^a Required to be worldwide unique and recommend to be EUI-64 concatenated with an 8 byte extension.</p> <p>^b The iSCSI name is a worldwide unique UTF-8 string no more than 224 bytes long, including the null character required to terminate the string.</p> <p>^c The Initiator Session Identifier (ISID) is a UTF-8 encoded hexadecimal representation of a non-zero six byte integer. The ISID is represented in no more than 12 bytes, not including null character termination if any.</p> <p>^d The Target Portal Group Tag (TPGT) is a UTF-8 encoded hexadecimal representation of a non-zero two byte integer. The TPGT is represented in no more than 4 bytes, not including null character termination if any.</p> <p>^e See IEEE Std P1212 for more information on the Discovery ID.</p>						

See table A.4 for a list of the size of the names for each SCSI transport protocol.

Table A.4 — Object name size for each SCSI transport protocol

Object	Name size ^a					
	SPI-5	FCP-2	SRP	iSCSI ^b	SBP-3	SAS SSP
Initiator device	not specified	not specified	not specified	224 bytes	not specified	8 bytes
Target device	not specified	not specified	not specified	224 bytes	not specified	8 bytes
Initiator port	not specified	8 bytes	16 bytes	245 bytes	8 bytes	not specified
Target port	not specified	8 bytes	16 bytes	245 bytes	11 bytes	not specified
Logical unit	Reported in the Device Identification VPD page (see SPC-3).					
^a Any SCSI transport protocol may support the SCSI name string format (see SPC-3), resulting in names with the sizes shown in the iSCSI column. ^b Including the terminating null character byte.						

See table A.5 for a list of the format of the names for each SCSI transport protocol.

Table A.5 — Object name format for each SCSI transport protocol

Object	Name format ^a					
	SPI-5	FCP-2	SRP	iSCSI	SBP-3	SAS SSP
Initiator device	not specified	not specified	not specified	SCSI name string format	not specified	NAA IEEE Registered format
Target device	not specified	not specified	not specified	SCSI name string format	not specified	NAA IEEE Registered format
Initiator port	not specified	Fibre Channel name_ identifier	EUI-64 8 byte extension ^c	iSCSI name ^b ",i,0x" Initiator Session Identifier ^d	EUI-64	not specified
Target port	not specified	Fibre Channel name_ identifier	EUI-64 8 byte extension ^c	iSCSI name ^b ",i,0x" Target Portal Group Tag ^e	EUI-64 Discovery ID ^f	not specified
Logical unit	Device Identification VPD page name (see SPC-3)					
<p>Key: means "concatenated with"</p> <p> ",i,0x" means a UTF-8 string composed of the following five characters: comma, lowercase i, comma, zero, and lowercase x.</p> <p> ",t,0x" means a UTF-8 string composed of the following five characters: comma, lowercase t, comma, zero, and lowercase x.</p>						
<p>^a In addition to the name formats shown in this table, any SCSI transport protocol may support the SCSI name string format (see SPC-3).</p> <p>^b The iSCSI name is a worldwide unique UTF-8 string no more than 224 bytes long, including the null character required to terminate the string.</p> <p>^c Required to be worldwide unique and recommend to be EUI-64 concatenated with an 8 byte extension.</p> <p>^d The Initiator Session Identifier (ISID) is a UTF-8 encoded hexadecimal representation of a non-zero six byte integer. The ISID is represented in no more than 12 bytes, not including null character termination if any.</p> <p>^e The Target Portal Group Tag (TPGT) is a UTF-8 encoded hexadecimal representation of a non-zero two byte integer. The TPGT is represented in no more than 4 bytes, not including null character termination if any.</p> <p>^f See IEEE Std P1212 for more information on the Discovery ID.</p>						

A.3 SCSI transport protocol acronyms and bibliography

A.3.1 EUI-64 (Extended Unique Identifier, a 64-bit globally unique identifier): The IEEE maintains a tutorial describing EUI-64 at <http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>.

A.3.2 FCP-2: SCSI Fibre Channel Protocol -2 (see 1.3).

A.3.3 IEEE Std P1212: Standard for a Control and Status Register (CSR) Architecture for Microcomputer Buses. See <http://www.ieee.org/>.

A.3.4 iSCSI: As of this writing, the most recently published iSCSI internet draft is: <http://www.ietf.org/internet-drafts/draft-ietf-ips-iscsi-20.txt>. Newer drafts may be identified at <http://http://www.ietf.org/html.charters/ips-charter.html>. The iSCSI internet draft is a standards track RFC specification.

A.3.5 NAA: Name Address Authority (see SPC-3).

A.3.6 SAS: Serial Attached SCSI (see 1.3).

A.3.7 SAS SSP: SAS (see A.3.6) Serial SCSI Protocol.

A.3.8 SBP-3: Serial Bus Protocol -3 (see 1.3).

A.3.9 SPI-5: SCSI Parallel Interface -5 (see 1.3).

A.3.10 SRP: SCSI RDMA Protocol (see 1.3).

A.3.11 UTF-8: See ISO/IEC 10646-1:2000, Information technology - Universal Multiple-Octet Coded Character Set (UCS) - Part 1: Architecture and Basic Multilingual Plane. See <http://www.iso.org/>.

Annex B

(informative)

Terminology mapping

The introduction of a model for SCSI devices with multiple ports resulted in changes in terminology between SAM and SAM-3 (see table B.1).

Table B.1 — SAM-3 to SAM terminology mapping

SAM-3 equivalent term	SAM term
initiator port identifier	initiator identifier
SCSI initiator port	initiator
SCSI port	port
SCSI port identifier	device identifier
SCSI port identifier	SCSI identifier
SCSI target port	target
target port identifier	target identifier