

T10/04-103r0

Voting Results on T10 Letter Ballot 04-102r0 on
Forwarding SAM-3 to First Public Review
Ballot closed: 2004/04/21 12:00 noon MDT

Organization	Name	S	Vote	Add'l	Info
Adaptec, Inc.	Tim Symons	P	Yes		
Agilent Technologies	Pat Thaler	P	Yes		
AMCC	Neil Wanamaker	P	Yes	Cmnts	
Amphenol Interconnect	Michael Wingard	P	Yes		
Brocade	Robert Snively	P	Yes	Cmnts	
Cisco Systems, Inc.	Claudio DeSanti	P	Yes		
Crossroads Systems, Inc.	Dexter Anderson	A	Yes	Cmnts	
Dallas Semiconductor	James A. Lott, Jr.	P	Yes		
Dell, Inc.	Kevin Marks	P	Yes		
EMC Corp.	David Black	A	Yes	Cmnts	
Emulex	Robert H. Nixon	P	Yes	Cmnts	
ENDL	Ralph O. Weber	P	Yes	Cmnts	
FCI	Douglas Wagner	P	Yes		
Fujitsu	Mike Fitzpatrick	P	Yes		
General Dynamics	Nathan Hastad	P	Yes		
Hewlett Packard Co.	Rob Elliott	P	No	Cmnts	
Hitachi Cable Manchester	Zane Daggett	P	Yes		
Hitachi Global Storage Tech.	Dan Colegrove	P	Yes		
IBM Corp.	George O. Penokie	P	No	Cmnts	
Intel Corp.	Robert Sheffield	P	Yes		
Iomega Corp.	David Hawks	P	Yes		
LSI Logic Corp.	John Lohmeyer	P	Yes		
Madison Cable Corp.	Ashlie Fan	P	Yes		
Maxtor Corp.	Mark Evans	P	Yes	Cmnts	
Microsoft Corp.	Emily Hill	P	Yes		
Molex Inc.	Jay Neer	P	Yes		
Panasonic Technologies, Inc	Terence J. Nelson	P	Yes		
Philips Electronics	William P. McFerrin	P	Yes		
Pivot3, Inc.	Bill Galloway	P	Yes		
QLogic Corp.	Skip Jones	P	Yes		
Quantum Corp.	Paul Entzel	P	Yes		
Seagate Technology	Gerald Houlder	P	Yes		
Storage Technology Corp.	Erich Oetting	P	Yes		
Sun Microsystems, Inc.	Charles Binford	A	Yes	Cmnts	
Texas Instruments	Paul D. Aloisi	P	Yes	Cmnts	
Toshiba	Hiroshi Suzuki	P	Yes		
UNISYS	Ron Mathews	P	Yes		
Veritas Software	Roger Cummings	P	Yes	Cmnts	

Ballot totals: (36:2:0:0=38)

36 Yes

2 No

0 Abstain

0 Organization(s) did not vote

38 Total voting organizations

4 Duplicate ballot(s) not counted

12 Ballot(s) included comments

This 2/3rds majority ballot passed.

36 Yes are more than half the membership eligible to vote minus abstentions
[greater than 19] AND

36 Yes are at least 26 (2/3rds of those voting, excluding abstentions [38])
AND

36 Yes are equal to or exceed a quorum [12]

Key:

P Voter is principal member

A Voter is alternate member
 Abs Abstain vote
 DNV Organization did not vote
 Cmnts Comments were included with ballot
 NoCmnts No comments were included with a vote that requires comments

Comments attached to Yes ballot from Neil Wanamaker of AMCC:

AMCC-01 (E) Page: 3.1.15 Location: 7

Problem Description:
 Second sentence nonsensical unless you already know how this works.

Suggested Solution:
 A CDB may either have a length up to 16 bytes determined by the group (bits 7:5 of the operation code) or a variable length from 12-260 bytes for operation code 7f.

AMCC-02 () Page: 3.1.20 Location: 8

Problem Description:
 conf -> status or class 2???

Suggested Solution:

AMCC-03 (E) Page: 3.1.21 Location: 8

Problem Description:
 This seems too restrictive - a rewind command that has not yet returned status is current by most definitions.

Suggested Solution:
 change "in the process of sending" to "has not yet sent"

AMCC-04 (E) Page: 3.1.56 Location: 10

Problem Description:
 "sent to a remote entity" - the data could be inbound.

Suggested Solution:
 change "sent to a remote entity" to "sent across the service delivery subsystem"

AMCC-05 () Page: 3.1.60 Location: 10

Problem Description:
 Logical units process task management functions

Suggested Solution:
 change "process commands" to "process commands and task management functions"

AMCC-06 (E) Page: 3.1.64, 4.8 Location: 10, 32

Problem Description:
 Appears inconsistent with Tables A.1, A.2; SBP appears to support 16-bit LUNs.

Suggested Solution:
 make consistent

AMCC-07 (E) Page: 3.1.79 Location: 11

Problem Description:

" usually according to the temporal order in which they were created."

Suggested Solution:

perhaps: "according to the order in which the tasks will be executed" (or remove).

AMCC-08 (T) Page: 4.7.1 Location: 27

Problem Description:

"Application clients are the sources of commands and task management functions." FCP-2 12.3.5 permits ABTS from a target. This also appears to duplicate the contents of the next paragraph. Note that this is not unique to FCP-2.

Suggested Solution:

Remove first of the two paragraphs; rephrase last paragraph to say "An application client task may originate a singlea"

AMCC-09 (T) Page: 4.7.1 Location: 27

Problem Description:

No service is described to process a received task management function (see previous comment).

Suggested Solution:

Describe a mechanism for an initiator to process received task management functions.

AMCC-10 (T) Page: 4.7.3 Location: 30

Problem Description:

Third from last paragraph "When the SCSI Target/Initiator devicea". See previous two comments.

Suggested Solution:

Modify wording to allow the task router to route task management functions even if acting as an initiator.

AMCC-11 (E) Page: 4.7.6 Location: 30

Problem Description:

"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." Per table 28, the remaining task management functions are all related to an I_T_L or I_T_L_Q nexus.

Suggested Solution:

Remove.

AMCC-12 (T) Page: 4.13.3 Location: 45

Problem Description:

"The REPORT LUNS commands (see SPC-3) shall be accepted by the logical unit with the logical unit number zero.." does not account for the REPORT LUNS well-known logical unit.

Suggested Solution:

Add "or, if the REPORT LUNS well-known logical unit is supported in the target and the command is addressed to the REPORT LUNS well-known logical unit, "

AMCC-13 (T) Page: 5.3.1, Task Set Full status Location: 63

Problem Description:

It appears that the "should" and "shall" in the first paragraph are reversed. If the expectation is that the initiator waits until a task completion before redriving an operation that received Task Set Full, the initiator should never receive a Task Set Full when there are no tasks from this initiator that could complete. Otherwise, there is no meaningful difference between Task Set Full and Busy.

Suggested Solution:

Swap "should" and "shall".

AMCC-14 (T) Page: 5.7.3, 5.9.7, Table 27 Location: 73, 82, 85

Problem Description:

5.7.3 suggests that if TAS is set to 0, a UA is generated; none of the conditions in 5.9.7 seems appropriate to this condition, nor do any of the entries in Table 27. It would appear that COMMANDS CLEARED BY ANOTHER INITIATOR would be appropriate in 5.9.7 and Table 27.

Suggested Solution:

As suggested.

AMCC-15 (T) Page: 5.7, 7.2 Location: 72-73, 90

Problem Description:

It is explicitly stated that the ABORT TASK does not change the state of reservations or mode parameters. If an ABORT TASK is issued for a MODE SELECT or PERSISTENT RESERVE OUT, is the state of the mode pages or reservation known? Is media unchanged on aborted writes?

Suggested Solution:

Clarify, or explicitly note that the device state is not known.

AMCC-16 (T) Page: 7.4 Location: 91

Problem Description:

It is not specified whether the issuance of a CLEAR ACA task management function is limited to the initiator that issued the command that created the ACA, and how this is affected by TST.

Suggested Solution:

Specify.

AMCC-17 (E) Page: Table a.2 Location: 109

Problem Description:

SPI-5 is referenced.

Suggested Solution:

remove

AMCC-18 (E) Page: 5.9.2.3.1, note 8 Location: 79

Problem Description:

the from SCSI initiator ports

Suggested Solution:

swap words

Comments attached to Yes ballot from Robert Snively of Brocade:

Brocade 001 (T) Page: ? Location: ?

Problem Description:

At first glance, the parameter lists for the execute command procedure call require auto-sense. It looks to me, again at first glance, as if even SPI-5 allows devices without auto-sense to operate. If this is true, then SAM-3 requires wording that includes saving sense information for REQUEST SENSE, probably as part of the definition of the execute command in clause 5.1 and perhaps other places.

Suggested Solution:

Consider explicitly indicating either that all modern SCSI protocols use auto-sense or including appropriate modifications to the text in 5.1 and perhaps other places.

Brocade 002 (E) Page: 10 Location: 3.1.65

Problem Description:

Media information is poorly defined here and used in only one place in the whole document. I believe the explanation should be expanded where it is used, in clause 5.2, second paragraph and removed from the glossary.

Suggested Solution:

In clause 5.2, "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." should be changed to "For all commands, if the logical unit detects an invalid parameter in the CDB, then the logical unit shall complete the command without executing any operations that will change any state in the device controlled by a SCSI command set."

Brocade 003 (E) Page: 26 Location: 4.6.3, last paragraph

Problem Description:

"about, or places" s/b "about, nor places"

Suggested Solution:

Make requested change

Brocade 004 (T) Page: 29 Location: 4.7.3

Problem Description:

The description b)B) appears to be exactly the same as a target device and an initiator device co-residing in a platform with no other connections between them. This function may be described, but it is NOT a SCSI target/initiator device.

Suggested Solution:

Modify the text to treat co-residing SCSI target devices and SCSI initiator devices as just another curious example, but not as one having the authority of a special type. That requires deletion of b)B) and separation of Figure 14 into a separate and explanatory (and perhaps unnecessary) clause.

Brocade 005 (E) Page: 31 Location: 4.7.8

Problem Description:

I believe that this text correctly allows multiple virtual N_Ports to be implemented as defined by FC-FS. However, it is possible to read this to make that impossible. Suggest review of text with this point in mind.

Suggested Solution:

Review 4.7.8 wording.

Brocade 006 (T) Page: 57 Location: 4.15

Problem Description:

The terminate data transfer request appears to me to be at a strange level in this explanation. I would have expected it to be a characteristic of the transport protocol called upon by such functions as task management. Where is there an example of its pure use at this level?

Suggested Solution:

This is basically a question to start with, but may require removal of this function.

Brocade 007 (T) Page: 59 Location: 5.1

Problem Description:

There is a fundamental problem in the definition of the Execute Command procedure call. One of the characteristics of channel architectures (as opposed to network architectures) is the presence of a pre-established Data-In Buffer Descriptor. This descriptor includes both length information and a handle or address that identifies a particular destination area. This is a key to the high performance "zero-copy" behavior of storage devices on channels. The Execute Command procedure call does not provide a complete Data-In Buffer Descriptor, but only a length. This Descriptor is required implicitly by all SCSI transports, but explicitly by SRP and other RDMA approaches to SCSI.

Note that this is not a direct requirement for all Data-Out Buffers, since the Data-Out Buffer is implicitly self describing, but something like it could be used for SRP and other RDMA transports.

Suggested Solution:

"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))" should be changed to read:

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

The description for Data-In Buffer Size should be rewritten as:

"Data-In Buffer Descriptor: A descriptor providing information about the length of the Data-In Buffer and the identifier of the Data-In Buffer location."

Alternatively, the Data-In Buffer Descriptor may be replaced with a Data-In Buffer Size and a Data-In Buffer Location Identifier, both required if a Data-In Buffer is used.

Brocade 008 (E) Page: 60 Location: 5.1

Problem Description:

"TASK COMPLETE, LINK ED COMMAND" s/b "TASK COMPLETE, LINKED COMMAND"

Suggested Solution:

Make requested change

Brocade 009 (T) Page: 65 Location: 5.4.2

Problem Description:

See Brocade 007.

Suggested Solution:

Make same changes to Execute Command request/confirmation as for Execute Command procedure call.

Brocade 010 (T) Page: 70 Location: 5.4.3.4

Problem Description:

Terminate Data Transfer Service does not appear to be a required function. It overlaps with the Task Management functions in a very uncertain way, terminating both data transfers for a particular command and all others. I believe that this should be handled instead by making sure that the text in clause 7, especially for ABORT TASK, include the option of terminating data transfers. If this is an artifact of the SAS state machine, it is not sufficiently general to be justified in SAM.

Suggested Solution:

Remove clause 5.4.3.4 and other references to Terminate Data Transfer service.

Include in clause 5.7.1 a definition of the functions that are realized when a Task is aborted. (This needs to be done anyway, since there is still some ambiguity about time relationships where an ABORT TASK is busy chasing a task through its entire request/response cycle and may catch up to it at any point in the cycle, leaving state that must also be cleared beyond the point in the cycle where the ABORT TASK becomes effective.)

The functions need to include:

- 1) Clear any remaining protocol state associated with the task.
- 2) Terminate any on-going data transfers associated with the task as possible.

Note that 2 allows the possibility that data transfers may continue for some short time after ABORT TASK has been received.

Brocade 011 (T) Page: 71 Location: 5.5

Problem Description:

It is implicit in the paragraph on deferred errors that deferred errors are cleared only by those events listed. In fact, deferred errors are reported by a subsequent completed command, but may be cleared by all kinds of actions, including ACA recovery operations.

Suggested Solution:

Allow the list to include dynamic Application Client sourced action as one mechanism for clearing deferred errors.

Brocade 012 (E) Page: 73 Location: 5.7.3

Problem Description:

"shall be an unit attention" s/b "shall be a unit attention"

Suggested Solution:

Make requested change

Brocade 013 (E) Page: 77 Location: 5.9.2

Problem Description:

This is a hanging paragraph. It looks like an editorial slip, since much of the text of 5.9.2 is redundant with the text of 5.9.2.1.

Suggested Solution:

Rewrite 5.9.2.1 to include any critical text of 5.9.2 not already included in 5.9.2.1 and delete the text of 5.9.2.

Brocade 014 (E) Page: 79 Location: 5.9.2.3.1

Problem Description:

"action) the from" s/b "action) from"

Suggested Solution:
Make requested change

Brocade 015 (E) Page: 81 Location: 5.9.4

Problem Description:
"an REQUEST" s/b "a REQUEST"

Suggested Solution:
Make requested change

Brocade 016 (T) Page: 82 Location: 5.9.6

Problem Description:
See Brocade 001. This strongly implies that SPI type protocols without autosense are not allowed under SAM-3, but this is not stated elsewhere, and SAM-3 does list the SPI series of documents in clause 1, implicitly supporting them.

Suggested Solution:
State explicitly that non-auto-sense behavior as defined in SPI (and I believe in many other SCSI command sets and some other protocols) is not supported by SAM-3 compliant devices. This should be done in Clause 1, either by eliminating such devices from the list or by explicitly disclaiming that support for each relevant document. A few words in 5.9.6 saying the same thing would be helpful.

Brocade 017 (E) Page: 89 Location: 7.1

Problem Description:
"shall only used by" s/b "shall only be used by"

Suggested Solution:
Make requested change

Brocade 018 (E) Page: 90 Location: 7.2

Problem Description:
The first sentence says the function shall be supported by all logical units. The third paragraph says "If the logical unit supports this function,"

Suggested Solution:
Delete "If the logical unit supports this function,"

Brocade 019 (T) Page: 92 Location: 7.8

Problem Description:
The text of this section implies that each SCSI transport protocol shall have a single mechanism for transmitting all Task Management requests. Several have different mechanisms for different commands. As one example, ABORT TASK is performed by a recovery abort sequence in FCP-2, but the remainder of the functions is achieved by a Task Management tag in the Command IU.

Suggested Solution:
Change "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." to read "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. The SCSI transport protocol standard may define more than one mechanism for implementing the Send Task Management Request transport protocol service, depending on the Task Management Request to be transported."

Brocade 020 (T) Page: 94 Location: 7.8

Problem Description:

The last paragraph allows task management functions to be requested and the response to be returned unrelated, requiring additional mechanisms to restrict the SCSI initiator port to a single pending task management request. I believe that all transport mechanisms should have a method to identify the responses and peg them to a particular request as a requirement. This is actually a separate question, since it implies that that because the tags of the effected tasks may be uncertain or many, that the task management response is somehow incapable of being identified.

Suggested Solution:

The last paragraph should read: "The specification of the SCSI transport protocol shall allow the Received Task Management Function Executed, confirming completion of the requested task, to be associated with the Task Management Request."

Brocade 021 (E) Page: 98 Location: 8.5.2

Problem Description:

"media). This" s/b "media) this

Suggested Solution:

Make requested change

Brocade 022 (T) Page: 102 Location: 8.8, Figure 39

Problem Description:

Note b should read "For ordered tasks, all older tasks have ended."

Suggested Solution:

Make requested change

Brocade 023 (T) Page: 104 Location: 8.9.2, Figure 40 and 41.

Problem Description:

The blocking boundary for task 3 should be in the same location as the blocking boundary for task 1. Head of Queue does not order with respect to subsequent tasks, but with respect to all simple and ordered tasks.

Suggested Solution:

Move blocking boundary for task 3 to the same location as for task 1 in snapshot 2.

The same change is required for figure 41.

Brocade 024 (T) Page: 104 Location: 8.9.1

Problem Description:

The last sentence implies that a task always enters the enabled task state after intervening barriers have been removed. In fact, it MAY enter the enabled task state, depending on the vendor specific behavior of the device.

Suggested Solution:

Change "enters the enabled" to "may enter the enabled".

Brocade 025 (E) Page: 110 Location: A.2, Table A.3

Problem Description:

There are some font problems in the FCP-2 column, Initiator Port row and Target Port row

Suggested Solution:

Correct font errors.

Comments attached to Yes ballot from Dexter Anderson of Crossroads Systems, Inc.:

Crossroads 1

Editorial

Pages 76&77 5.9.1.1 & 5.9.2

The text in these two sections is similar.

I recommend 5.9.1.1 be rewritten to more closely match 5.9.2.

Comments attached to Yes ballot from David Black of EMC Corp.:

EMC letter ballot comments on SAM-3r13. All of these comments are editorial.

[1] 4.9.2 LUN 0 Address

The last sentence of this section refers to the peripheral device address method.

This is the first mention of this address method, and it is a forward reference -

please add a note that this method is specified in 4.9.6.

[2] 4.9.3 Single level logical unit number structure.

The descriptions of Tables 1 and 2 contain forward references to 4.14:

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

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

Aside from the forward references, these are incorrect because 4.14 describes

a model, not a format. Please remove the references to 4.14 from these two sentences and add a sentence at the start of 4.9.3 saying that the logical unit number formats defined in this section are for a single level subset of the hierarchical model for dependent logical units described in 4.14.

[3] 4.9.7 Flat space addressing method and 4.9.8 Extended logical unit addressing

Both of these sections appear to be missing the equivalent of the first paragraph

of 4.9.5 and 4.9.6 requiring that the received command be sent to the logical

unit. That requirement should be present in all four sections, or none of them.

[4] 4.13.2 SCSI devices with multiple ports

In the following sentence:

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.

Both "that meet product requirements" and "performed by the product" are excess verbiage - please remove them.

[4] 4.13.2 SCSI devices with multiple ports

The meaning of the following sentence is unclear:

The structures and views of SCSI devices are asymmetric for SCSI target ports and SCSI initiator ports.

Depending on what was intended, either replace "asymmetric" with "different"

or remove the sentence if it is a restatement of the previous sentence.

[5] 4.14 Model for dependent logical units

The description of Figure 25 in terms of whether or not additional SCSI target

devices or SCSI domains can be added is peculiar. The 3 dots notation used in the figure generally denotes omission of things that actually exist for the purpose of simplification. The description should be rewritten to better

correspond to the figure.

[6] 5.3.1 Status codes

The following text appears to be over-broad:

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).

The "shall" could be interpreted as requiring CONDITION MET to be returned in a large number of situations where it is inappropriate to do so. I suggest adding an initial sentence to say that CONDITION MET is only used with specific commands (e.g., SBC PRE-FETCH commands) that allow it, and those commands specify its usage. The above sentence would then begin with "For such commands, this status shall be returned ..."

[7] 5.3.2 Status precedence

Note 5 is troublesome. Where are the unit attention conditions defined/specified, 5.9.7? Something other than "previous versions of this standard" seems to be needed here.

[8] 5.4.3.1 Introduction

Near the end of the section:

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 phrase "command data" is potentially confusing, delete the word "command",

or rephrase.

[9] A.3.4 iSCSI

The reference will be RFC 3720 in the near future.

Comments attached to Yes ballot from Robert H. Nixon of

Emulex:

Emulex comment number

E for Editorial, T for Technical

physical page number

subclause

issue

proposed resolution

001

E

pp21

1.3

The paragraph above Figure 2 says "Figure 2 shows the relationship..." The paragraph below Figure 2 says "Figure 2 is not intended to imply a relationship..." Figure 2 is so vague as to be meaningless, even without this assistance.

Delete Figure 2 and replace the first three paragraphs of 1.3 with "The SCSI standards family comprises standards in the following five functional areas:"

002

T

21

3.1

As the arbiter of the SCSI model, it would be useful for SAM to point out that Interconnects need not be specific to SCSI.

After the second sentence describing Interconnects, add a new sentence "Interconnect standards may not be specific to SCSI transport."

003

T

A.2

127

In table A.2, the initiator port identifier size for iSCSI is shown as 246. According to footnotes b and c in table A.3, this looks like it should be 241

(including the trailing null). Also, it is the maximum size, not the fixed size.

In table A.2, change the iSCSI initiator port identifier size to 241 bytes (maximum) and change the iSCSI target port identifier size to 233 bytes(maximum).

004

T

128

A.2

In table A.3, it is stated that the identifier for an iSCSI Target port uses the infix ",i,0x". I believe it should be ",t,0x"

In table A.3, in the identifier for an iSCSI Target port change the infix ",i,0x" to ",t,0x"

005

T

128

A.2

Footnote b in table A.3 could be interpreted to require a trailing null in the middle of the identifiers to which it applies.

In table A.3, change footnote b to "The iSCSI name is a worldwide unique UTF-8 string no more than 223 bytes long. As it is used only with a suffix in this table, it does not include a null character to terminate the string."

006

T

A.2

129

In table A.4, the initiator and target port name sizes for iSCSI are shown as 246. According to feetnotes b, d and e in table A.5, this looks like initiator should be 241 and the target should be 233 (each including the trailing null). Also, they are the maximum sizes, not the fixed sizes.

In table A.4, change the iSCSI initiator port name size to 241 bytes(maximum) and change the iSCSI target port name size to 233 bytes(maximum).

007

T

130

A.2

In table A.5, it is stated that the name for an iSCSI Target port uses the infix ",i,0x". I believe it should be ",t,0x"

In table A.5, in the name for an iSCSI Target port change the infix ",i,0x" to
",t,0x"

008

T

130

A.2

Footnote b in table A.5 could be interpreted to require a trailing null in the middle of the names to which it applies.

In table A.5, change footnote b to "The iSCSI name is a worldwide unique UTF-8 string no more than 223 bytes long. As it is used only with a suffix in this table, it does not include a null character to terminate the string."

Comments attached to Yes ballot from Ralph O. Weber of
ENDL:

ENDL 1 Technical

pg 68, 5.4.3.1, Editor's Note 1

Remove Editor's Note 1. Make no other changes.

Comments attached to No ballot from Rob Elliott of
Hewlett Packard Co.:

PDF Page xiv
2.7 Revision 6

"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."
SBP-3 still only supports a 2 byte LUN field. Some of these changes may need to be undone. A sentence could be added in 4.9 stating "Some transport protocols only support a single level of the 8 byte LUN structure."

HPQ #2

PDF Page 2

1.2 Requirements precedence

After "command descriptor block" add "(CDB)"

This section precedes the acronym section, so it might be too early to just use CDB alone.

HPQ #3

PDF Page 7

3.1.16 command standard

Change "command standard" to "command set standard" which is used more often

HPQ #4

PDF Page 7

3.1.3 application client

Change "commands" to "commands and task management functions."

HPQ #5

PDF Page 7

3.1.15 command descriptor block (CDB)

Add "See 5.2 and SPC-3."

HPQ #6

PDF Page 8

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."

(for Hugh Curley)

This doesn't clearly state that the domain also *contains* the service delivery subsystem, which figure 9 in 4.5 does indicate.

Reword as:

"An I/O system consisting of a set of SCSI devices and a service delivery subsystem, where the SCSI devices interact with one another by means of the service delivery subsystem."

HPQ #7

PDF Page 8

3.1.21 current task

Delete "Each SCSI transport protocol standard should define the SCSI transport protocol specific conditions under which a task is considered a current task."

This is unnecessary. Protocols already have to define Send Data-In and Receive Data-Out protocol services. This should suffice.

HPQ #8

PDF Page 9

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."

Delete this definition. Ssee main comment on this topic in 4.6.1.

HPQ #9

PDF Page 10

3.1.57 implicit head of queue

"commands wherein the specified commands"

Reword in terms of tasks, since tasks have task attributes, not commands (until linked are obsoleted).

HPQ #10

PDF Page 11

3.1.85 request-confirmation transaction

This term is not used anywhere. The header "5.4.2 Execute Command request/confirmation SCSI transport protocol services" is the closest

HPQ #11

PDF Page 12

3.1.110 sense key

Change "A field" to "The SENSE KEY field"

HPQ #12

PDF Page 12

3.1.96 SCSI port

If request/response terminology is intended, leave as is.
If protocol service is intended, change "requests and responses are routed." to "requests, indications, responses, and confirmations are routed."

HPQ #13

PDF Page 12

3.1.95 SCSI initiator port

If protocol service terminology is intended, leave as is.
If generic request/response terminology is intended, change "requests and confirmations are routed." to "requests and responses are routed."

HPQ #14

PDF Page 12

3.1.100 SCSI target port

"through which indications and responses are routed."

If this is using the 4 protocol service terms, it also services requests and confirmations for data transfers (in the opposite direction).

Change to

"through which device server requests, indications, responses, and confirmations are routed."

If this is using generic request-response terms, change to "through which requests and responses are routed."

(see comment on 3.1.95)

HPQ #15

PDF Page 13

3.1.116 signal

Is this term still needed?

"(n) A detectable asynchronous event possibly accompanied by descriptive data and parameters.

(v) The act of generating such an event."

HPQ #16

PDF Page 13

3.1.117 standard INQUIRY data

After "an INQUIRY command." add "with the EVPD bit set to zero (see SPC-3)."

HPQ #17

PDF Page 13

3.1.124 task

Add "See 4.11."

HPQ #18

PDF Page 14

3.1.135 unlinked command

After "A command having the LINK bit set to zero in the CDB CONTROL byte."
add "not preceded by a command that had the
LINK but set to one in the CDB CONTROL byte."

Otherwise this conflicts with the definition of linked command in 3.1.59,
which includes a command with LINK=0 as the last one in
the sequence.

HPQ #19

PDF Page 19

4.1 Introduction (RC)

item b

Whole sentence is too long and difficult to follow.

Suggestions:

'Identify areas for developing standards and provide a common reference for
maintaining consistency among related standards. In
this way, independent implementers may work productive and independently .'

HPQ #20

PDF Page 20

4.2 SCSI distributed service model (RC)

In the last paragraph, the way they explain 'client-server relationships
not being symmetrical' is not clear enough. I'm not sure
what they are trying to achieve in that paragraph. I think they are mixing
2 concepts.

Maybe there should be a split into 2 paragraphs and an addition:

'.A server may only respond to such requests. In other words, a client may
not have the ability to behave like a server and vice
versa.

[New paragraph] The client requests .'

HPQ #21

PDF Page 21

4.3 The SCSI client-server model

Change "command standards" to "command set standards"
(see comment in 3.1.16)

HPQ #22

PDF Page 21

4.3 The SCSI client-server model

"command completion response is sent" is wrong. A task will result in lots
of SCSI Command Completes. Change to "task
complete response".

HPQ #23

PDF Page 23

4.4 SCSI structural model (RC)

The figure should have an explanation on what the symbols mean.

For example, the diagonal lines mean 0 or more objects. If they are not
present then the box means 1 or more objects. And so on
with the rest of the boxes in the diagram.

HPQ #24

PDF Page 24

4.6.1 The service delivery subsystem object

Get rid of the "interconnect subsystem" level, which adds little value.

The "Service delivery subsystem" term can stand on its own.

Miscellaneous comments in 4.6.1 and 3.1 implement this.

(reprise of an EMC comment on SAM-2 in 02-155)

HPQ #25

PDF Page 24

4.6.1 The service delivery subsystem object

Delete "and is composed of an interconnect subsystem (see figure 10)."
leaving just:

"The service delivery subsystem connects SCSI ports (see 3.1.96)"
See main comment on this topic in 4.6.1.

HPQ #26

PDF Page 24

4.6.1 The service delivery subsystem object

Delete "Figure 10 – Service delivery subsystem model"

See main comment on this topic in 4.6.1.

HPQ #27

PDF Page 24

4.5 SCSI domain

Change "command" to "command or task management function"

HPQ #28

PDF Page 24

4.5 SCSI domain

Change "commands" to "commands and task management functions"

HPQ #29

PDF Page 24

4.5 SCSI domain

Change "commands" to "commands and task management functions"

HPQ #30

PDF Page 24

4.5 SCSI domain

Change "commands" to "commands and task management functions"

HPQ #31

PDF Page 24

4.5 SCSI domain

Change "commands" to "commands and task management functions"

HPQ #32

PDF Page 24

4.5 SCSI domain (RC)

1st sentence maybe include some more words

'A SCSI domain is composed of at least one SCSI devices, which includes at least one SCSI target port .'

HPQ #33

PDF Page 25

4.6.1 The service delivery subsystem object

Delete "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."
(see main comment on this topic in 4.6.1)

HPQ #34

PDF Page 25

4.6.1 Service delivery subsystem object

"Considered received" by whom"? The sender doesn't know until the protocol service returns back a confirmation.

Changed c) to "Considered received by the receiver" and a) to "Considered sent by the sender"

HPQ #35

PDF Page 25

4.6.3 Request/Response ordering (RC)

2nd sentence contains an example that is confusing. It may be re-worded differently.

'. and may take action based on the nature and sequence of SCSI target device responses. An example of how/when

non-ordering could go wrong is if the SCSI initiator device aborts .' [parenthesis have been removed]

HPQ #36

PDF Page 25

4.6.1 Service delivery subsystem object

"requests, responses, and data"

If generic request/response terminology is intended, change to "requests and responses"

(if this paragraph is kept at all)

HPQ #37

PDF Page 25

4.6.3 Request/Response ordering

Are data transfers included in "responses"? Does this mean the generic request/response type of response, or the protocol service response?

(same comment in 7.2)

HPQ #38

PDF Page 26

4.6.3 Request/Response ordering

The first sentence speaks very generally: "The SCSI architecture model assumes in-order delivery to be a property of the service delivery subsystem."

The third sentence then makes a statement that belies the first: "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."

If it truly makes no assumption about different SCSI initiator ports, then the first sentence is too strong. Restrict the first sentence to one initiator port.

"The SCSI architecture model assumes in-order delivery _between tasks and task management functions from a single SCSI initiator port_ to be a property of the service delivery subsystem." (another comment asks that "of requests and responses" be added after "delivery" too)

HPQ #39

PDF Page 26

4.6.3 Request/Response ordering

The first sentence "The SCSI architecture model assumes in-order delivery to be a property of the service delivery subsystem." doesn't specify in-order delivery of _what_.

The third sentence mentions "tasks or task management functions," but is mentioning that there is NO assumption for them for different SCSI initiator ports. It is silent about ordering for the same initiator port. 4.8 mentions that this standard does not require in-order delivery or processing of task management functions (but isn't clear if it "assumes" it).

The preceding paragraphs describe requests and responses in general. The protocol service requests defined later in this standard are:

Send Command Complete, Send Data-In, Receive Data-Out, Send Task Management Request, Terminate Data Transfer

The responses are:

Send Command Complete, Task Management Function Executed

Are those all "assumed" to be ordered, even the data transfer requests?

Assuming so...

Reword the first sentence as: "The SCSI architecture model assumes in-order delivery _of requests and responses_ to be a property of the service delivery subsystem."

Reword the third sentence as:

"This standard makes no assumption about, nor places any requirement on, the ordering of requests or responses between tasks or task management functions received from different SCSI initiator ports." Then, to clarify what "assumes" means and doesn't mean, add:
"Although written assuming in-order delivery of requests and responses, the SCSI architecture model does not require in-order delivery of requests and responses."

HPQ #40

PDF Page 26

4.6.3 Request/response ordering

"This standard makes no assumption about, or places any requirement on the ordering..."

or should be nor and a command should be added.

Change to:

"This standard makes no assumption about, nor places any requirement on, the ordering...."

HPQ #41

PDF Page 28

4.7.2 SCSI target device

Change "be accessed using the logical unit number zero." to "be LUN 0."

HPQ #42

PDF Page 30

4.7.4 SCSI port identifier

Delete "The SCSI port identifier is equivalent to SCSI identifier." There is no need for the imprecise "SCSI identifier" term any more.

HPQ #43

PDF Page 30

4.7.3 SCSI target/initiator device

Change "be accessed using the logical unit number zero." to "be LUN 0."

HPQ #44

PDF Page 30

4.7.6 SCSI task router (MB)

"The task router routes tasks..."

This statement contradicts 5.5, Task and Command Lifetimes, which states, 'The device server shall create a task upon receiving a SCSI Command Received indication' See also 5.8.1 and 5.8.2. The Task Router actually routes commands, data, and status associated with (or bound to) a task.

HPQ #45

PDF Page 30

4.7.6 SCSI task router (MB)

A device server creates a task [5.5]. A device server can only exist inside of a logical unit [4.8]. Hence the only way for a task to be sent to a logical unit is for a device server to create the task and then send it to some a logical unit. I know of no other location in SCSI that supports this view of tasks. Replace 'task' with 'SCSI Command Received indication'.

HPQ #46

PDF Page 30

4.7.5 Relative port identifier (RC)

It is not clear if one relative port identifier is used for all ports in the SCSI device; or if a group of ports in the SCSI device will share the same relative port identifier; or if there is a different relative port identifier for every SCSI port in the SCSI device. I found it confusing, so it needs clarification.

HPQ #47

PDF Page 32

4.8 Logical units

When multi-level LUNs are used, is the nested logical unit really a "logical unit within the scope of a SCSI target device"? In SCC, a nested LUN often represents LUN 0 on a disk drive on a bus behind the RAID controller. Is the RAID controller expected to intercept the INQUIRY VPD data from that LU and change the reported target port name (association=2) and the reported target port identifier (association=1)?

The current proposal for bridges in MSC plans to use multi-level LUNs to access devices behind a bridge. The bridge will not be expected to do so (that's one of the points of the proposal).

Change "within the scope" phrases to something like "accessible via the"

HPQ #48

PDF Page 32

4.8 Logical units

"A logical unit number is a field"

In this context, the logical unit number is not really a field, it's just a value.

Change to "A logical unit number contains 64 bits..."

HPQ #49

PDF Page 32

4.8 Logical units

"may require that a logical unit include a logical unit name" implies that it is not always present. If that is the case, the b) item above needs to change from "one or more logical unit names" to "zero or more".

Or better yet, require that a logical unit name be included in all logical units except well-known logical units. Don't make that a transport protocol option. Leave b) alone.

HPQ #50

PDF Page 32

4.8 Logical units

Change

"a) A logical unit number;"

to

"a) One or more logical unit numbers: " or "a) Logical unit number(s)" because the B) that follows allows more than one.

HPQ #51

PDF Page 32

4.8 Logical units

Delete "per logical unit;" which is redundant with the "A logical unit contains" introduction

HPQ #52

PDF Page 32

4.8 Logical units

Delete "per logical unit;" which is redundant with the "A logical unit contains" introduction

HPQ #53

PDF Page 32

4.8 Logical units

Change "b) one or more logical unit names" to "b) one or more logical unit names if the logical unit is not a well-known logical unit"

HPQ #54

PDF Page 32

4.8 Logical units

Change "may contain" to "contains"
 Since zero or more is contained, there's no need for "may"

HPQ #55

PDF Page 32

4.8 Logical units

Since well-known logical units are not allowed to have logical unit names, the "logical unit name" box should be a zero or more box with diagonal lines.

HPQ #56

PDF Page 32

4.8 Logical units

Add "There is one device server per logical unit." to match the sentence in the next paragraph about the task manager.

4.8 Logical unit

Change "carries out" to "processes"

HPQ #58

PDF Page 32

4.8 Logical units

Move this paragraph into 4.6.3 Request/response ordering.

"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."

HPQ #59

PDF Page 33

4.9.3 Single level logical unit number structure

If this text is kept (other comments suggest removing it)...

Change "in 4.14" to "in 4.9.6". 4.14 is the dependent logical unit model; the address method 00h format is defined in 4.9.6.

HPQ #60

PDF Page 33

4.9.2 LUN 0 address

"All SCSI devices shall accept LUN 0 as a valid address." needs to be clarified with respect to 5.9.4 Incorrect logical unit selection. If PQ=001 at LUN 0, it is not going to support most commands (but has to support REPORT LUNS, REQUEST SENSE, and INQUIRY).

HPQ #61

PDF Page 33

4.9.2 LUN 0 address

"To address the LUN 0 of a SCSI device the peripheral device address method shall be used." implies that LUN 0 or other LUN numbers might be accessed via more than one address method. The shall hints that there are choices that are not allowed.

However, LUN 0 by definition has an address method of 00b. Some other LUN with another address method is not LUN 0; it is a different LUN (e.g. LUN 40000000_00000000 is different from LUN 00000000_00000000).

Change to "LUN 0 is addressed with the peripheral device address method."

HPQ #62

PDF Page 33

4.9.3 is confusing, because the rules it contains also apply to the last level of the eight byte logical unit number field as well. Merge this section into the eight-byte LUN sections: 4.9.4 (peripheral device)

for the format for LUNs 255 and below, and 4.9.7 (flat space) for LUNs 16383 and below)

HPQ #63

PDF Page 33

4.9 Logical unit numbers

There is terminology conflict between LUN referring to the 8-byte structure, and a variety of subfields also called LUNs (in whole or in part):

Table 1, 2 - SINGLE LEVEL LUN (8 bits or 14 bits)

Table 1, 2 - Null second level LUN (16 bits)

Table 1, 2 - Null third level LUN (16 bits)

Table 1, 2 - Null fourth level LUN (16 bits)

Table 7 - LUN field (5 bits)

Table 8 - TARGET/LUN (8 bits)

Table 9 - LUN (14 bits)

Since the rest of SCSI uses LUN to mean the 8-byte structure, all these subfields should be renamed. Table 4 calls each of the 2 byte structures:

FIRST LEVEL ADDRESSING

SECOND LEVEL ADDRESSING

THIRD LEVEL ADDRESSING

FOURTH LEVEL ADDRESSING

Other comments suggest removing tables 1 and 2, folding their information into the peripheral device and flat space sections.

That would leave these needing clarification:

Table 7 - LUN field (5 bits) - change to PARTIAL LUN field

Table 8 - TARGET/LUN (8 bits) - change to TARGET OR PARTIAL LUN field

Table 9 - LUN (14 bits) - change to PARTIAL LUN field

HPQ #64

PDF Page 34

4.9.3 Single level logical unit number structure

Change "in 4.14" to "in 4.9.7"

HPQ #65

PDF Page 34

4.9.3 Single level LUN format

"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."

Reword as:

Devices with 256 or fewer logical units should assign them LUNs using the peripheral device addressing method.

and place in 4.9.4 (see other comments about merging 4.9.3 into 4.9.4+)

HPQ #66

PDF Page 34

4.9.3 Single level LUN format

"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."

Reword as:

Devices with 16384 or fewer logical units should assign them LUNs using the flat space addressing method.

HPQ #67

PDF Page 34

4.9.3 Single level LUN structure

"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."

This implies that the same LU might be accessed with LUNs 00000000_000000nn or 40000000_000000nn. Really those are two different LUs.

Delete this paragraph.

HPQ #68

PDF Page 35

4.9.4 Eight byte logical unit number structure

Figure 16

Change Byte to byte

HPQ #69

PDF Page 35

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."

"Four levels of devices under a single device" is the same as "5 levels of SCSI devices"; I think the latter would be clearer if levels of devices are being described. Figure 25 in 4.14 calls itself a 3 level example; to match the terminology in this sentence, it would be more appropriate to call it a 2 level example.

The logical unit address method 10b sneaks in an extra level, since it specifies separate bus number, target, and LUN values. (It must be the last in the list).

Example:

LUN going into first level SCSI device: 1,2,3,4 (addr method 01b, bus number A target B, LUN C)

LUN going into second level SCSI device: 2,3,4

LUN going into third level SCSI device: 3,4

LUN going into fourth level SCSI device: 4

The fourth level device parses the 10b format and relays the command out bus number A to target B, addressing it to LUN C. That is received by a fifth level SCSI device.

Reword like: "The eight byte LUN structure contains four levels of addressing fields."

HPQ #70

PDF Page 36

4.9.4 Eight byte logical unit number structure

Table 4

Change Byte to byte

HPQ #71

PDF Page 36

4.9.4 Eight byte LUN structure

Table 4 - Eight byte LUN structure

Delete each (MSB) and (LSB) in this table (8 total) since each pair of bytes has sub-structures

HPQ #72

PDF Page 36

4.9.4 Eight byte LUN format

"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."

If a logical unit is pointed to by level X, however, the device at level X-1 (the previous level) must have the ability to route the request to level X (e.g., be an SCC-2 device)

HPQ #73

PDF Page 36

4.9.4 Eight byte LUN structure

Table 5 - format of addressing fields

Remove (MSB) and (LSB) since the ADDRESS METHOD SPECIFIC field has substructure

HPQ #74

PDF Page 37

4.9.x

All the "addressing methods" should use mixed case. It is very difficult to parse the logical unit not specified extended address method.

Globally change terms to:

logical unit addressing method -> Logical Unit addressing method

peripheral device addressing method -> Peripheral Device addressing method

flat space addressing method -> Flat Space addressing method

extended logical unit addressing method -> Extended Logical Unit addressing method

well known logical unit extended address -> Well Known Logical Unit extended address

logical unit not specified extended address method -> Logical Unit Not Specified extended address

HPQ #75

PDF Page 37

4.9.5 Logical unit addressing method

Table 7

Change the two cells containing 1 0 into one cell containing "ADDRESS METHOD (10b)"

HPQ #76

PDF Page 37

4.9.4 Eight byte logical unit number structure

Sort the addressing methods in table 6 and sections 4.9.5 through 4.9.8 in this order:

00b Peripheral device

01b Flat space

10b Logical unit

11b Extended logical unit

This is numerically sorted and more closely matches their level of use. The current order is not even alphabetical.

HPQ #77

PDF Page 37

4.9.5 LU addressing method

Change "ADDRESS METHOD SPECIFIC field" to "address field". The table has 2 bits more than just the ADDRESS METHOD SPECIFIC field contents.

HPQ #78

PDF Page 37

4.9.5 LU addressing method

The LUN field doesn't seem to be usable except when the LU addressing format appears last in the chain.

What would it mean to have a LUN of:

bytes 0-1: address method 10b, target=XX, bus number=YY, LUN=ZZ

bytes 2-3: anything but 00h

bytes 4-5: 00h

bytes 6-7: 00h

Assume that LUN is sent to device A. The target and bus number fields make sense; device A chooses output bus YY and sends the command to target XX on that bus. The LUN field on that bus, however, cannot be ZZ. It has to be the shifted version:

bytes 0-1: anything but 00h

bytes 2-5: 000000h

There's no way to include the original "LUN" value of ZZ on that output bus or otherwise use it for relaying.

Add a rule that the LUN field is ignored unless this addressing field is the last in the 8 byte LUN. Or, that if this addressing field appears, none of the subsequent address fields in the 8 byte LUN are used.

HPQ #79

PDF Page 37

4.9.4 Eight byte LUN structure

Table 6 Address Method field values

Add a note to 01b that "This address method is called the Volume set addressing method by SCC-2"

HPQ #80

PDF Page 37

4.9.5 Logical unit address method

4.9.6 Peripheral device address method

(and, if changed by another comment, 4.9.7 Flat space 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."

The SCSI device also has to relay task management functions. Those cannot be terminated with CHECK CONDITION status.

What if it does not choose to relay one of them?

Add "or task management function" to the paragraph and the following note. Describe the response if it is not relayed.

HPQ #81

PDF Page 37

4.9.5 Logical unit addressing method

"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)."

There are no SCSI transport protocols covered by SAM-3 that define target port identifiers that fit into this 6 bit field, so they all have to be mapped representations. Delete the note and change the sentence to:

"The TARGET field contains a mapped representation of a target port identifier."

HPQ #82

PDF Page 38

4.9.6 Peripheral device addressing method

"The BUS IDENTIFIER field may use the same value encoding as the BUS NUMBER field (see 4.9.5)."

This is a 6 bit field, but the referred-to Bus Number field is 3 bits. They can't completely use the same encoding.

HPQ #83

PDF Page 38

4.9.6 Peripheral device addressing method

Table 8

Change the two cells containing 0 0 into one cell containing "ADDRESS METHOD (00b)"

HPQ #84

PDF Page 38

4.9.6 Peripheral device addressing method

Change "ADDRESS METHOD SPECIFIC field" to "address field". The table has 2 bits more than just the ADDRESS METHOD SPECIFIC field contents.

HPQ #85

PDF Page 38

4.9.7 Flat space addressing method

Change:

"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."

to:

"If the flat space 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."

or, if this address method is considered special because it just has a LUN field:

"If the flat space addressing method is selected the SCSI device shall relay the received command to the current level logical unit specified by the LUN field."

This paragraph was just copied from SCC-2, where it did match the wording in the other addressing formats. Since it's in SAM-3, it should match the wording in the other SAM-3 paragraphs describing addressing formats.

HPQ #86

PDF Page 38

4.9.7 Flat space addressing method

Delete "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)."

which doesn't seem to belong here. It may have made sense in SCC-2 where it was pulled from.

HPQ #87

PDF Page 38

4.9.6 Peripheral device addressing

Table 8 and throughout the section

Change the "TARGET/LUN" field to the "TARGET OR LUN" field.

The current name sounds like it's a single value that contains a combination of target and LUN values somehow mapped into one byte.

Really, it is based on BUS IDENTIFIER:

if BUS IDENTIFIER is 00h, contains a (small) LUN value

if BUS IDENTIFIER is not 00h, contains a (mapped) target port identifier

HPQ #88

PDF Page 38

4.9.6 Peripheral device addressing

Delete "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)."
There is no SCSI fan or cache command set, so why would those ever need LUNs? These examples don't make sense. The rules from 4.9.3 should be merged into this section instead.

HPQ #89

PDF Page 38

4.9.6 Peripheral device addressing method

Delete "within or joined to the current level SCSI device."

HPQ #90

PDF Page 38

4.9.6 Peripheral device addressing method

"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."
There are no SCSI transport protocols covered by SAM-3 that define target port identifiers that fit into this 8 bit field, so they all have to be mapped representations. Change the sentence to:
"The TARGET field contains a mapped representation of a target port identifier."

HPQ #91

PDF Page 38

4.9.6 Peripheral device addressing method

Change "the logical unit with the logical unit number zero" to "LUN 0"

HPQ #92

PDF Page 39

4.9.7 Flat space addressing method

Table 9

Change the two cells containing 0 1 into one cell containing "ADDRESS METHOD (01b)"

HPQ #93

PDF Page 39

4.9.8 Extended logical unit addressing

Table 10

Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #94

PDF Page 39

4.9.7 Flat space addressing method

Change "ADDRESS METHOD SPECIFIC field" to "address field". The table has 2 bits more than just the ADDRESS METHOD SPECIFIC field contents.

HPQ #95

PDF Page 39

4.9.8 Extended logical unit addressing

Change "ADDRESS METHOD SPECIFIC field" to "address field". The table has 2 bits more than just the ADDRESS METHOD SPECIFIC field contents.

HPQ #96

PDF Page 39

4.9.8 Extended logical unit addressing

Change "addressing" to "addressing method" to match 4.9.5, 4.9.6, and 4.9.7

HPQ #97

PDF Page 39

4.9.8 Extended logical unit addressing

Change "EXTENDED ADDRESS METHOD SPECIFIC field" to "extended address field". The length refers to 2 bits more than just the EXTENDED ADDRESS METHOD SPECIFIC field contents.

HPQ #98

PDF Page 40

4.9.8 Extended logical unit addressing

Table 12

Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #99

PDF Page 40

4.9.8 Extended logical unit addressing

Table 13

Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #100

PDF Page 40

4.9.8 Extended logical unit addressing

Table 12

Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #101

PDF Page 40

4.9.8 Extended logical unit addressing

Table 12

Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #102

PDF Page 40

4.9.8 Extended logical unit addressing

Table 16

Change "Code" to "Code(s)" to match the length column. Check the font of "Code".

HPQ #103

PDF Page 40

4.9.8 Extended logical unit addressing

Table 13, 14, 15

Delete (MSB) and (LSB). There are no multibyte EXTENDED ADDRESS SPECIFIC fields defined yet. When they are defined, they may or may not have substructures and the MSB/LSB labels might be inappropriate.

HPQ #104

PDF Page 41

4.9.9 Well known logical unit addressing
and

4.9.10 Logical unit not specified addressing
should be subsections of 4.9.8 Extended logical unit addressing
since they are both of type 11b

HPQ #105

PDF Page 41

4.9.9 Well known logical unit addressing

Table 17 - Well known LU extended address format

"Well known logical unit (1h)"

The usual format for this is

"EXTENDED ADDRESS METHOD (1h)"

HPQ #106
PDF Page 41
4.9.9 Logical unit not specified addressing
Table 17 - Logical unit not specified addressing
"Logical unit not specified (Fh)"
The usual format for this is
"EXTENDED ADDRESS METHOD (Fh)"

HPQ #107
PDF Page 41
4.9.9 Well known logical unit addressing
Change "well" to "the well"

HPQ #108
PDF Page 41
4.9.10 Logical unit not specified addressing
If an 8-byte LUN has an extended address method LUN in the second, third, or fourth position, but the LENGTH field extends past the end of the LUN, what should happen?
The only format currently defined that is not 2 bytes is the logical unit not specified format. If 2 bytes of FFFFh appears in the LUN field in any position, are the prior SCSI devices allowed to relay it and zero fill at the end? That means the last device could see an incoming LUN like FFFF0000_00000000.
It seems like that should be treated as a logical unit not specified LUN. So, table 18 should just show a 2-byte addressing field instead of an 8-byte field. If the extended addressing method is FFh and the length is 11b, it shall be treated as logical unit not specified, regardless of the remaining bytes in the LUN. Perhaps even treat any multi-byte extended address format that exceeds past the end of of the LUN as a logical unit not specified.

HPQ #109
PDF Page 41
4.9.10 Logical unit not specified addressing
Table 18
Delete "method" from the table 18 header. It's not in table 17 or preceding tables.

HPQ #110
PDF Page 41
4.9.9 Well known logical unit addressing
Change "EXTENDED ADDRESS METHOD SPECIFIC field" to "address field". The table has 2 bits more than just the EXTENDED ADDRESS METHOD SPECIFIC field contents.

HPQ #111
PDF Page 41
4.9.9 Well known logical unit addressing
Table 17
Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #112
PDF Page 41
4.9.10 Logical unit not specified addressing
Table 18
Change the two cells containing 1 1 into one cell containing "ADDRESS METHOD (11b)"

HPQ #113
PDF Page 42
4.10 Well known logical units
Change invalid to incorrect

HPQ #114
PDF Page 42

4.10 Well known logical units

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

into a NOTE, since this is not the place for such a rule. It's really a note for item b) explaining why target device names is plural.

HPQ #115
PDF Page 42

4.11 Tasks and task tags

Add c) Optionally, a task priority (see 8.7)

HPQ #116
PDF Page 43

4.12 The nexus object

Define when a nexus is created and destroyed. Different for I_T, I_T_L, and I_T_L_Q. I_T is protocol specific. I_T_L_Q is pretty well defined.

SAM-2 HP comment (see 02-155) was:

This does not describe when each of the nexus objects comes into existence, and when it is destroyed (issue raised in 02-078r1).

The following is suggested - I_T nexus object is instantiated upon the first successful instantiation of an I_T_L_x nexus object as indicated by the SCSI protocol layer interactions. The I_T nexus object is destroyed on receiving the "I_T Nexus loss" notification from the SCSI protocol (Rob Elliott's 02-134r0). The I_T_L nexus object is instantiated when the first valid task to the LU is received and accepted (i.e. the task enters the Dormant state) and destroyed when the corresponding I_T nexus object is destroyed. The I_T_L_Q nexus object is instantiated when the corresponding I_T_L nexus object is already instantiated (thus exists) and when a task with a tag Q is issued on the nexus. The I_T_L_Q nexus object is destroyed on the conclusion of the said task, or when the I_T_L nexus object is destroyed.

Editor's Note: It is very difficult to describe when an I_T nexus is instantiated in a manner that is compatible with both parallel SCSI and Fibre Channel. It may be possible to address this issue in SAM-3.

HPQ #117
PDF Page 44

4.13.2 SCSI devices with multiple ports

Delete "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)."

This is not true. The obsolete "SCSI identifier" term now means SCSI port identifier. A port might have multiple association=1 identifiers in VPD data, but it isn't treated as multiple ports because of that.

HPQ #118
PDF Page 45

4.13.3 Multiple port target SCSI device structure

Change "SCSI port name or identifier values" to "logical unit names".

The port names/identifiers don't help discovery of the same logical unit through multiple target ports; they help discover the same port through multiple logical units. This sentence is backwards.

(see similar comment in 4.14.5)

HPQ #119
PDF Page 45

4.13.3 Multiple port target SCSI device structure
Change "ports, however, communications" to "ports. However, communications"
HPQ #120
PDF Page 45

4.13.3 Multiple port target SCSI device structure
Change "target SCSI device" to "SCSI target device"

HPQ #121
PDF Page 45

4.13.3 Multiple port target SCSI device structure
"Each SCSI target port shall accept commands sent to LUN 0 and the task router shall route them to a device server for processing." is not completely true. If LUN 0 has PQ=001b, only REQUEST SENSE, INQUIRY, and REPORT LUNS must be accepted. The rest can result in CC/ILLEGAL REQUEST/LOGICAL UNIT NOT SUPPORTED"

HPQ #122
PDF Page 45

4.13.3 Multiple port target SCSI device structure
Change "the logical unit number zero" to "LUN 0"

HPQ #123
PDF Page 45

4.13.3 Multiple port target SCSI device structure (MB)
Link 'inactive' to the asymmetric logical unit access states defined in SPC3, 5.8.4.

HPQ #124
PDF Page 46

4.13.4 Multiple port initiator SCSI device structure
Change "initiator SCSI device" to "SCSI initiator device"

HPQ #125
PDF Page 46

4.13.4 Multiple port initiator SCSI device structure
Change "as if no such mechanisms exist" to "on a port, not device, basis"

HPQ #126
PDF Page 46

4.13.4 Multiple port initiator SCSI device structure
Figure 19
The application clients should be shown together, all having access to any initiator port they select.
Execute Command and the Send SCSI Command protocol service each take an I_T_L_Q nexus argument, which implies that any application client can select any initiator port (any I) to service its request. If they couldn't, it'd be a T_L_Q argument.
(this comment is appears for Figure 20, 22, and 23.)

HPQ #127
PDF Page 47

4.13.5 Multiple target/initiator ports
Change "SCSI port name or identifier values" to "logical unit names".
The port names/identifiers don't help discovery of the same logical unit through multiple target ports; they help discover the same port through multiple logical units. This sentence is backwards.
(see similar comment in 4.13.3)

HPQ #128
PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure
Change "target/initiator SCSI device" to "SCSI target/initiator device"

HPQ #129

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure

After "Figure 20 shows the structure of a SCSI target/initiator device with multiple SCSI target/initiator ports."

add "A SCSI target/initiator device may also contain SCSI target ports and SCSI initiator ports." so this picture doesn't imply that only target/initiator ports are allowed

HPQ #130

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure

Delete "a" or change to "a single"

HPQ #131

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure

Figure 20

Show the application clients with access to all initiator ports

HPQ #132

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure

Change "the logical unit number zero" to "LUN 0"

HPQ #133

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure

"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."

See comment in 4.13.3 on the same sentence.

HPQ #134

PDF Page 47

4.13.5 Multiple port target/initiator SCSI device structure (MB)

Link 'inactive' to the asymmetric logical unit access states defined in SPC3, 5.8.4.

HPQ #135

PDF Page 48

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

Change domains to SCSI domains

HPQ #136

PDF Page 48

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

Figure 21

This almost implies that the SCSI domain contains 4 service delivery subsystems. It really contains 1. A would be better.

HPQ #137

PDF Page 48

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

Since this is describing figure 21 where there is only one SCSI target device, it should be singular.

The "access via target port identifier" concept is confusing; use the i.e. instead.

change "...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." to "...the SCSI

target device are accessible via multiple SCSI target ports and map the configuration of the SCSI target device."

HPQ #138
PDF Page 49

4.13.6 SCSI initiator device view of a multiple port SCSI target device
Figure 22
Show two clouds. It looks like each domain has two service delivery
subsystems.

HPQ #139
PDF Page 49

4.13.6 SCSI initiator device view of a multiple port SCSI target device
Change "multiple ports to "multiple SCSI initiator ports and multiple SCSI
target ports"

HPQ #140
PDF Page 49

4.13.6 SCSI initiator device view of a multiple port SCSI target device
Figure 23
Show a cloud so it doesn't look like the SCSI domain has 4 service delivery
subsystems.

HPQ #141
PDF Page 49

4.13.6 SCSI initiator device view of a multiple port SCSI target device
Figure 23
Show the application clients with access to all initiator ports

HPQ #142
PDF Page 49

4.13.6 SCSI initiator device view of a multiple port SCSI target device
Figure 22
Show the application clients with access to all initiator ports

HPQ #143
PDF Page 50

4.13.5 Multiple port target/initiator SCSI device structure
"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."
Not true.
VPD page 83h logical unit names provide this info, since a logical unit
must be in only one target device. The target device name
directly provides this information, as well. The first sentence hints at
application client knowledge.
Since the application client is part of the device, the second sentence is
untrue. It might be true that the SCSI initiator port (in
some protocols) cannot make a distinction, but does that matter?
Delete the sentence.

HPQ #144
PDF Page 51

4.14 Model for dependent logical units
This is only used by Controller (SCC) devices today (and possibly Bridge
(MSC) devices in the future). Add a statement to that
effect.

HPQ #145
PDF Page 51

4.14 Model for dependent logical units
Change "the logical unit number zero" to "LUN 0"

HPQ #146
PDF Page 60

5.1 Execute Command procedure call (CM)
"(e.g., sense data, to determine the state of the buffer contents)."
Closing ")" should be after "sense data" not at end of sentence.

HPQ #147
PDF Page 61
5.2 CDB

After "sense key of ILLEGAL REQUEST" add "and an additional sense code of
INVALID FIELD IN CDB"

HPQ #148
PDF Page 61
5.2 CDB (CM)
"media information."
Presumably more than just "media information" should not change; e.g. mode
parameters, log parameters, etc.
Something along the lines of "The commanded action shall not be carried
out" (better words needed!)

HPQ #149
PDF Page 62
5.2 CDB
Add a paragraph break before "If the" so BUSY looks like TASK SET FULL and
RESERVATION CONFLICT.

HPQ #150
PDF Page 62
5.3.1 Status codes and global
Change "an unit" to "a unit"
<http://www.stclaresoxfordonline.fsworld.co.uk/pages/langprac/articles-indef.htm>
shows "an university", which has the same sound,
is incorrect.

HPQ #151
PDF Page 63
5.3.1 Status codes
The return of ACA ACTIVE needs to exempt PR OUT with a PREEMPT AND ABORT
service action (as described in 5.9.2.3.1):
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.

HPQ #152
PDF Page 63
5.2 CDB
Delete "or an element of a logical unit" since element reservations are now
obsolete in SPC-3

HPQ #153
PDF Page 63
5.2 CDB
Delete the period in "SPC-3)."

HPQ #154
PDF Page 63
5.3.1 Status codes (CM)
TASK SET FULL
This status shall be implemented by all logical units.
Seems strange that this is the only one explicitly called out as "shall be
implemented by all logical units" - that must be true for
others too (GOOD, CHECK)?

HPQ #155

PDF Page 64

5.3.2 Status precedence (CM)

There are other unit attention conditions which would seem pretty important; I'm thinking of:

LVD Transceivers changed state

Microcode has been changed

Are these lower precedence because they're not listed explicitly, or are they overlooked?

HPQ #156

PDF Page 64

5.3.2 Status precedence (CM)

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

It's not clear to me how this paragraph relates to the preceding one. "any level of precedence" could mean:

"after the above highest-precedence conditions"

OR

"wherever you like; could take precedence over the above highest-precedence conditions"

Is it deliberately vague, or not explicit by accident?

HPQ #157

PDF Page 65

5.4.2 Execute Command request/confirmation SCSI transport protocol services Delete "Support for the SCSI Command Received indication and Send Command Complete response by a SCSI transport protocol standard is optional."

Why wouldn't they be mandatory? The only SCSI implementation that might not have them would be a software stack. In that case, it's not a "SCSI transport protocol".

HPQ #158

PDF Page 65

5.4.2 Execute Command (JL)

"All SCSI I/O systems shall implement these SCSI transport protocols as defined in the applicable SCSI transport protocol specification."

should these applicable specifications be listed?

HPQ #159

PDF Page 65

5.4.2 Execute Command (JL)

shorten sentence to make easier to read?

HPQ #160

PDF Page 65

5.4.2 Execute Command (JL)

Data-In Buffer Size:

Data-Out Buffer:

Data-Out Buffer Size:

Command Reference

Number (CRN):

Task Priority:

First Burst Enabled:

Should these also be described as 'if present' as is in Send Command Complete (p66) for Sense Data?

HPQ #161

PDF Page 66

5.4.2 Execute Command (JL)

Status & Sense Data Length incorrectly ordered to prototype.

HPQ #162

PDF Page 66

5.4.2 Execute Command (JL)

Command Reference

Number (CRN):

Task Priority:

First Burst Enabled:

Should these also be described as 'if present' as is in Send Command Complete (p66) for Sense Data?

HPQ #163

PDF Page 66

5.4.2 Execute Command (JL)

Sense Data:

Sense Data Length:

Should Send Data Length also be described as 'if present' like Sense Data?

HPQ #164

PDF Page 67

5.4.3.1 Introduction

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

To avoid any implication that the Data-In Buffer and Data-Out Buffer might overlap, split this into two parts, one for Data-In Buffer and one for the Data-Out Buffer. Replace "Application Client Buffer Offset & Size" in each of those parts with Data-In Buffer Offset & Size and Data-Out Buffer Offset & Size.

HPQ #165

PDF Page 67

5.4.3.1 Introduction

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

Change "Byte Count Requested by Device Server" to "Request Byte Count" to match the name used later

HPQ #166

PDF Page 68

5.4.3.1 Introduction

Editor's Note 1

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

Proposal 03-002 had proposed changing each shall to should. However, every transport protocol is required to support one byte resolution for at least the last data frame of a command transferred in each direction. It is common for transport protocols to place restrictions on intermediate frames, which might correspond to Send Data In or Receive Data Out procedure calls, which should also be recognized by the text.

Change to:

All SCSI transport protocol standards shall define support for a resolution of one byte for the Data-In Buffer Size argument and the Data-Out Buffer Size argument.

SCSI transport protocol standards may define restrictions on the resolution of the Data-In Buffer Offset argument and the Data-Out Buffer Offset. [e.g. SAS requires they always be 4-byte aligned]

SCSI transport protocol standards may define restrictions on the resolution of the Request Byte count argument for all but the last call to Send Data-In and all but the last call to Receive Data-Out for a command. They shall support a resolution of one byte for the last call to Send Data-In and the last call to Receive Data Out for a command.

HPQ #167

PDF Page 68

5.4.3.1 Introduction

Change "Byte Count Requested by Device Server" to "Request Byte Count" two times on this page.

The latter is the name used in the data transfer protocol services.

HPQ #168

PDF Page 68

5.4.3.1 Introduction

"Application Client Buffer Size"

This is not an argument to a data transfer protocol service; it is used in Send SCSI Command/SCSI Command Received instead.

It has two names there - Data In Buffer Size and Data-Out Buffer Size - since bidirectional commands use both buffers.

Change "Application Client Buffer Size" to "Data-In Buffer Size or Data-Out Buffer size"

Change "The total number of bytes in the application client's buffer (Data-In or Data-Out)." to "The total number of bytes in the application client's buffer (Data-In or Data-Out), as specified by the application client in the Send SCSI Command protocol service and indicated to the device server by the SCSI Command Received protocol service"

HPQ #169

PDF Page 68

5.4.3.1 Introduction

Change "the combination of Application Client Buffer Size minus the Application Client Buffer Offset." to

"the combination of Data-In Buffer Size or Data-Out Buffer Size minus the Application Client Buffer Offset."

HPQ #170

PDF Page 68

5.4.3.1 Introduction (JL)

'Monotonically increasing' - use of obscure word? Monotonic - adj. Math. (of a function or quantity) varying in such a way that it either never decreases or never increases. Can we just say, 'increasing'?

HPQ #171

PDF Page 69

5.4.3.2 Data-In delivery service

Change "Application Client Buffer Offset" to "Data-In Buffer Offset" twice in this section so Data-In and Data-Out transfers can be concurrent and are not confused

HPQ #172

PDF Page 69

5.4.3.3 Data-Out delivery service

Change "Application Client Buffer Offset" to "Data-Out Buffer Offset" twice in this section so Data-In and Data-Out transfers can be concurrent and are not confused

HPQ #173

PDF Page 71

5.5 Task and command lifetimes (MB)

Clarify the reference to 'task'. To which task does this sentence refer, the device server task or the application client task? Since the sentence contains 'shall', I initially read 'task' as the device server task. However a service response of SERVICE DELIVERY OR TARGET FAILURE may mean that the device server did not create a task and hence cannot end it. This inconsistency makes me wonder if the sentence refers to the application client task.

HPQ #174

PDF Page 71

5.5 Task and command lifetimes (MB)

Terrible grammar. Change to 'shall end.'

HPQ #175

PDF Page 73

5.7.1 Mechanisms that cause tasks to be aborted

Delete "with or without establishing an ACA condition" or put in parenthesis. Since both cases are covered, it's just a parenthetical aside rather than part of the rule itself.

HPQ #176

PDF Page 73

5.7.3 When ... aborts task from another

"If the TAS bit is set to zero, the method of notification shall be an unit attention condition."

This appears to contradict SPC-3 clause 7.4.6 TAS bit, which says "A task aborted status (TAS) bit set to zero specifies that aborted tasks shall be terminated by the device server without any response to the application client."

"without any response" seems to preclude a unit attention.

Change SPC-3 to "the device server shall abort tasks without returning status." or otherwise make SAM-3 and SPC-3 more consistent.

HPQ #177

PDF Page 73

5.7.3 When ... aborts task from other ports

"The additional sense code set for the unit attention condition depends on the action that caused the task(s) to be aborted."

Enumerate the additional sense codes for each of the reasons in 5.7.1.

- a) CLEAR TASK SET yields COMMANDS CLEARED BY ANOTHER INITIATOR
- b) CC QErr=1 yields COMMANDS CLEARED BY ANOTHER INITIATOR
- c) PR OUT/PREEMPT AND ABORT yields COMMANDS CLEARED BY ANOTHER INITIATOR
- d) LOGICAL UNIT RESET yields BUS DEVICE RESET FUNCTION OCCURRED (see 6.2)

HPQ #178

PDF Page 79

5.9.2.3.1 Commands permitted from non-faulted... (MD)

Note 8

Delete "the"

HPQ #179

PDF Page 81

5.9.4 Incorrect logical unit selection

Change an to a

HPQ #180

PDF Page 81

5.9.4 Incorrect logical unit selection

After "terminated" add "by the SCSI target device"

HPQ #181

PDF Page 81

5.9.4 Incorrect logical unit selection

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

HPQ #182

PDF Page 81

5.9.4 Incorrect logical unit selection

"Any command except REQUEST SENSE or INQUIRY:" is not complete. REPORT LUNS has some special rules too. LUN 0 is supposed to support REPORT LUNS even if its PQ is 001b.

HPQ #183

PDF Page 81

5.9.4 Incorrect logical unit selection

Please tie the phrases like "does not support the logical unit" in all the A) B) entries to the PQ values returned by INQUIRY for that logical unit.

Either:

A) The logical unit returns a PQ of 011b in the Standard INQUIRY data (see SPC-3) (i.e., the SCSI target device does not support the logical unit)

or

A) The SCSI target device does not support the logical unit (i.e., the logical unit returns a PQ of 011b in the Standard INQUIRY data (see SPC-3);

HPQ #184

PDF Page 81

5.9.3 Overlapped commands

"that command" is not the command that was just received with the overlapping tag, nor is it the command/task already using that tag.

Change to "for that I_T_L_Q nexus."

HPQ #185

PDF Page 81

5.9.3 Overlapped command

Delete "10 Some logical units may not detect an overlapped command until after the CDB has been received."

This is a remnant of parallel SCSI where the tag was sent in a discrete step before the CDB was sent. In modern protocols it's in the same frame as the CDB.

HPQ #186

PDF Page 81

5.9.3 Overlapped commands

"A task manager that detects an overlapped command..."

Both FCP and SAS define task tags that have an I_T scope, not an I_T_L scope. If a command arrives with a tag in use in logical unit N, must only tasks from that initiator port in logical unit N aborted, or can tasks from that initiator port in all logical units be aborted?

HPQ #187

PDF Page 81

5.9.3 Overlapped commands

Add a cross reference to 5.1 to pick up the description of linked overlapped commands

HPQ #188

PDF Page 82

5.9.6 Sense data

UA_INTLCK_CTRL should be smallcaps.

HPQ #189

PDF Page 82

5.9.6 Sense data

After "other conditions" add "(e.g., with the REQUEST SENSE command (see SPC-3))"

HPQ #190

PDF Page 82

5.9.6 Sense data

After "field" add "in the Control mode page (see SPC-3)".

HPQ #191
PDF Page 88
6.4 Event notification (JL)
shorten sentence to make easier to read?

HPQ #192
PDF Page 90
7.3 ABORT TASK SET
Change "SCSI initiator port" to "I_T nexus".
Because:
1) The previous sentence said "abort all tasks...that were created by the SCSI initiator port routed through the SCSI target port"
2) the argument is I_T_L

HPQ #193
PDF Page 90
7.3 ABORT TASK SET
Change "SCSI initiator ports" to "I_T nexuses".
Because:
1) The previous sentence said "abort all tasks...that were created by the SCSI initiator port routed through the SCSI target port"
2) the argument is I_T_L

HPQ #194
PDF Page 90
7.2 ABORT TASK
Are data transfers included in "responses"? Does this mean the generic request/response type of response, or the protocol service response?
(same comment in 4.6.3)

HPQ #195
PDF Page 90
7.2 ABORT TASK
Delete "to the SCSI initiator port."

HPQ #196
PDF Page 92
7.8 Task management SCSI transport protocol services
After "Argument encoding the task management function to be performed."
add:
"The encoding between the application client and initiator port and between the target port and task manager is outside the scope of this standard, and between the SCSI initiator port and SCSI target port is protocol-specific."

HPQ #197
PDF Page 96
8.1 Intro to task set mgmt
"with a status other than GOOD."
doesn't account for CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET, which are also decent statuses.
Change to "with CHECK CONDITION status."
Transport protocol specific errors don't cause any of the other status values.

HPQ #198
PDF Page 96
4.3 Implicit head of queue
Change "command standard" to "command set standard"
(see comment in 3.1.16)

HPQ #199

PDF Page 96

8.1 Intro to task set management (MB)

Change 'controls application clients' to 'controls that application clients'.

HPQ #200

PDF Page 96

8.1 Intro to task set management (MB)

Change 'that task' to 'that cause the task'.

HPQ #201

PDF Page 96

8.1 Intro to task set management (MB)

The second a)b)c) list implies that a logical unit can perform work for a task that is not in any task set. The sentence, 'the requirements for task set management only apply ...' exempts tasks outside of all task sets from everything in clause 8. Hence the standard says nothing about the order in which the logical unit executes the task versus other tasks in its task set(s). The standard should state that the logical unit shall, when the device server creates a task, either place the task into a task set or end the task.

HPQ #202

PDF Page 98

8.5.1.2 Suspended information (MB)

Change 'is required to' to 'shall'.

HPQ #203

PDF Page 98

8.5.1.1 Task state nomenclature (MB)

Use of 'although' makes this sentence a fragment. Either remove 'although' or merge this sentence with the next one.

HPQ #204

PDF Page 98

8.5.2 Enabled task state (MB)

"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."

This sentence is too obscure. I think it states that the states a task passes through prior to entering enables task state shall not affect the subsequent behaviour of the task's command(s). However I can't be sure given the convoluted sentence structure.

HPQ #205

PDF Page 109

Annex A

Delete the SPI-5 columns in all the tables in annex A; SPI-5 goes with SAM-2 not SAM-3.

HPQ #206

PDF Page 113

A.3.4 iSCSI

Before completing SAM-3, see if iSCSI has been published as an RFC yet.

HPQ #207

PDF Page 113

A.3.4 iSCSI

Delete extra http://

Comments attached to No ballot from George O. Penokie of

IBM Corp.:

IBM-001

PDF pg 11, pg xi, Revision Information

The revision information needs to be removed before letter ballot

IBM-002

PDF pg 18, pg xviii, Introduction

The change bars should be removed.

IBM-003

PDF pg 22, pg 4, 1.3 SCSI standards family

There is no point in this list of standards. It is never 100% correct. It should be deleted from SAM-3.

IBM-004

PDF pg 25, pg 7, 3.1 Definitions

When you have a reference at the end of the definition you use the notation

<< text description (see xxx). >> in other case you use the notation << text

description. See xxx. >> I prefer the first but whichever is used it should be

used consistently throughout the definitions section.

IBM-005 Technical

PDF pg 25, pg 7, 3.1.15 command descriptor block (CDB):

The statement << variable length of between 12 and 260 bytes. >> does not seem

correct. The smallest variable length should be GT 16. If its smaller that 16

then it not really a variable length CDB.

IBM-006

PDF pg 28, pg 10, 3.1.65 media information:

The statement << non-volatile (retained through a power cycle) and >>

should

be << non-volatile (i.e., retained through a power cycle) and >> .

IBM-007

PDF pg 31, pg 13, 3.1.125 task tag:

The statement << 64 bits that is a component of an I_T_L_Q nexus. >> should be

<< 64 bits that is the Q component of an I_T_L_Q nexus. >>

IBM-008

PDF pg 39, pg 21, 4.3 The SCSI client-server model, 1st paragraph under figure 8

The statement << There is one application client task for each pending command, series of linked commands, or task management request. >> seems to be

a duplicate of the statement in the same paragraph that states << 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. >> and therefore should be deleted.

IBM-009 Technical

PDF pg 44, pg 26, 4.7 SCSI devices, 2nd paragraph

The statement << A SCSI target device contains at least one SCSI target port

and is capable ...>> should be << A SCSI target device contains at least one

SCSI target port, at least one logical unit, and is capable ...>> as a target

device without a logical unit is of no use in this model.

IBM-010 Technical

PDF pg 44, pg 26, 4.7 SCSI devices, 2nd paragraph

The statement << A SCSI target/initiator device contains at least one SCSI target/initiator port and is capable ... >> should be << A SCSI target/initiator device contains at least one SCSI target/initiator port,

at

least one logical unit, and is capable ... >> as a target device without a logical unit is of no use in this model.

IBM-011

PDF pg 44, pg 26, 4.7 SCSI devices, 2nd paragraph

The statement << a SCSI domain needs to contain >> should be << A SCSI domain

shall contain >> or << A SCSI domain is required to contain >>

IBM-012 Technical

PDF pg 46, pg 28, 4.7.2 SCSI target device, Last paragraph

The statement << shall be accessed using the logical unit number zero. >> should be changed to <<shall be accessed using the logical unit number zero or

a Report LUNs well-known logical unit. >> to allow for a target device to have

either a LUN 0 or a W-LUN.

IBM-013 Technical

PDF pg 48, pg 30, 4.7.3 SCSI target/initiator device, 2nd to last paragraph

The statement << shall be accessed using the logical unit number zero. >> should be changed to <<shall be accessed using the logical unit number zero or

a Report LUNs well-known logical unit. >> to allow for a target device to have

either a LUN 0 or a W-LUN.

IBM-014

PDF pg 50, pg 32, 4.8 Logical units, 1st paragraph after figure 15

The statement << dependent logical units in its composition, >> should have a

reference to the dependent logical unit section.

IBM-015

PDF pg 50, pg 32, 4.8 Logical units, 1st paragraph after figure 15

The statement << Otherwise, the logical unit numbers should >> is not exact and should be changed to << If there are no dependent logical units within the

scope of the SCSI target device, the the logical unit numbers should >>.

IBM-016

PDF pg 51, pg 33, 4.9.2 LUN 0 address

The title of section 4.9.2 << LUN 0 address >> should be changed to << Minimum

LUN addressing requirements >>.

IBM-017 Technical

PDF pg 51, pg 33, 4.9.2 LUN 0 address, 1st paragraph

The statement << All SCSI devices shall accept LUN 0 as a valid address. >> does not address the W-LUN option and should be changed to << All SCSI devices

shall accept LUN 0 or a well known logical unit as a valid address >>.

IBM-018 Technical

PDF pg 51, pg 33, 4.9.2 LUN 0 address, 1st paragraph

The statement << model the LUN 0 shall be the logical unit that >> does not address the W-LUN option and should be changed to << model the LUN 0 or a well

known logical unit shall be the logical unit that >>

IBM-019

PDF pg 51, pg 33, 4.9.3 Single level logical unit number structure, 1st paragraph

The statement << 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. >> is redundant with the statement in section 4.9.1 and should be deleted.

IBM-020

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure, 2nd paragraph under table 2

The statement << logical units should be numbered 255 and below. >> does not

clearly state that the address method should be 00b. Change to << logical units should be numbered 255 and below using the 00b address method. >>

IBM-021

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure, 3rd paragraph under table 2

The statement << logical units should be numbered 16 383 and below. >> does not clearly state that the address method should be 01b. Change to << logical

units should be numbered 16 383 and below using the 01b address method. >>

IBM-022

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure, 3rd paragraph under table 2

The statement << greater than 255 shall have >> should be changed to << greater than 255 and less than 16 384 shall have >>

IBM-023 Technical

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure, 3rd paragraph under table 2

This paragraph is a restatement of the two paragraphs directly above it. I think those two should be deleted.

IBM-024 Technical

PDF pg 56, pg 38, 4.9.6 Peripheral device addressing method, Last paragraph

The statement << 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). >> should be changed to << The SCSI device

located

within the current level may be addressed by a BUS IDENTIFIER field and a TARGET/LUN field of all zeros, also known as LUN 0 (see 4.9.2). >> as this is

no longer the only way to find out information about a SCSI device. The alternative is W-LUN.

IBM-025 Technical

PDF pg 63, pg 45, 4.13.3 Multiple port target SCSI device structure, 1st paragraph after figure 18

The statement << Each SCSI target port shall accept commands sent to LUN 0 and

the >> does not address the W-LUN option and should be changed to << Each SCSI

target port shall accept commands sent to LUN 0 or a well known logical unit

and the >>

IBM-026 Technical

PDF pg 63, pg 45, 4.13.3 Multiple port target SCSI device structure, 1st paragraph after figure 18

The statement <<The REPORT LUNS commands (see SPC-3) shall be accepted by the

logical unit with the logical unit number zero from >> does not address the

W-LUN option and should be changed to << The REPORT LUNS commands (see SPC-3)

shall be accepted by the logical unit with the logical unit number zero or the

report LUNs well known logical unit from >>

IBM-027 Technical

PDF pg 65, pg 47, 4.13.5 Multiple port target/initiator SCSI device structure,

1st paragraph after figure 20

The statement << Each SCSI target/initiator port shall accept commands sent to

LUN 0 and the >> does not address the W-LUN option and should be changed to

<<

Each SCSI target/initiator port shall accept commands sent to LUN 0 or a well

known logical unit and the >>

IBM-028 Technical

PDF pg 65, pg 47, 4.13.5 Multiple port target/initiator SCSI device structure,

1st paragraph after figure 20

The statement <<The REPORT LUNS commands (see SPC-3) shall be accepted by the

logical unit with the logical unit number zero from >> does not address the

W-LUN option and should be changed to << The REPORT LUNS commands (see SPC-3)

shall be accepted by the logical unit with the logical unit number zero or the

report LUNs well known logical unit from >>

IBM-029

PDF pg 69, pg 51, 4.14 Model for dependent logical units, figure 24

The dependent logical unit should be labeled << dependent logical unit >>.

IBM-030 Technical

PDF pg 70, pg 52, 4.14 Model for dependent logical units, 2nd paragraph after

figure 25

The statement << All addressable entities may default to vendor specific values or may be defined by an application client >> does not consider W-LUNs

and should be changed to << All addressable entities, except well known logical units, may default to vendor specific values or may be defined by an

application client >>

IBM-031

PDF pg 71, pg 53, 4.15 The SCSI model for distributed communications

The SAL, and STPL are already defined in the definitions section. So why redefined them here. They should be deleted.

IBM-032

PDF pg 71, pg 53, 4.15 The SCSI model for distributed communications
The interconnect layer is not defined in the definitions section but should be. Move it there and delete it from this section.

IBM-033

PDF pg 72, pg 54, 4.15 The SCSI model for distributed communications
The following terms are already defined in the definitions section therefore those definitions should be deleted from this section: SCSI transport protocol
service request, SCSI transport protocol service indication, SCSI transport protocol service response, and SCSI transport protocol service confirmation.

IBM-034

PDF pg 78, pg 60, 5.1 The Execute Command procedure call, Data-In Buffer description
The statement << 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. >> should be << As a result the contents of the buffer are not valid unless the command completes with a status of GOOD, CONDITION MET, INTERMEDIATE, or INTERMEDIATE-CONDITION MET. >>. There should be no assumptions in a standard over than it works perfectly.

IBM-035

PDF pg 79, pg 61, 5.2 Command descriptor block (CDB), Last paragraph
The statement << Bit 1 provides an obsolete way to request interrupts between

linked commands. >> should be deleted as the function is obsolete so no further statement need be stated.

IBM-036 Technical

PDF pg 81, pg 63, 5.3.1 Status codes, ACA Active

Add into the list the following: << If the TASK_ONLY bit is set to one and a

task with the ACA attribute is received from the faulted initiator port; >>

This will be in a upcoming proposal.

IBM-037

PDF pg 86, pg 68, 5.4.3.1 Introduction, Editor's Note

I don't have an answer to this but if none comes forward then it should not be

changed. In any case the note needs to be deleted.

IBM-038

PDF pg 89, pg 71, 5.5 Task and command lifetimes, 4th paragraph

The statement << The application client assumes that the task exists and maintains an application client task to interact with the task from the time

>> should be << The application client shall maintain an application client task to interact with the task from the time >>.

IBM-039

PDF pg 89, pg 71, 5.5 Task and command lifetimes, item A

The statement << COMMANDS CLEARED BY ANOTHER INITIATOR (if in reference to the

task set containing the task); >> should be << COMMANDS CLEARED BY ANOTHER INITIATOR if in reference to the task set containing the task; >

IBM-040

PDF pg 95, pg 77, 5.9.1.3 Aborting other tasks when CHECK CONDITION status is

returned without establishing an ACA, 1st paragraph

The statement << The TST (task set type) Control mode page field specifies >>

should be << The TST field specifies >> as the fact that it is in the control

mode page is already stated in this paragraph and there is no reason the give

the long version of the field name as that is done no where else.

IBM-041

PDF pg 95, pg 77, 5.9.1.3 Aborting other tasks when CHECK CONDITION status is

returned without establishing an ACA, 1st paragraph

The statement <<The QERR (queue error management) Control mode page field specifies how >> should be << The QERR field specifies how >> as the fact that it is in the control mode page is already stated in this paragraph and there is no reason the give the long version of the field name as that is done

no where else.

IBM-042 Technical

PDF pg 95, pg 77, 5.9.2 Auto contingent allegiance (ACA), 2nd paragraph

Change the statement << 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). >>

to << While the ACA condition is in effect and the TASK_ONLY bit is set to zero, 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). >> This will be in a upcoming proposal.

IBM-043

PDF pg 95, pg 77, 5.9.2.1 Establishing an ACA, 2nd paragraph

The statement << The TST (task set type) Control mode page field specifies >>

should be << The TST field specifies >> as the fact that it is in the control

mode page is already stated in this paragraph and there is no reason the give

the long version of the field name as that is done no where else.

IBM-044

PDF pg 96, pg 78, 5.9.2.1 Establishing an ACA, 2nd paragraph

The statement <<The QERR (queue error management) Control mode page field specifies how >> should be << The QERR field specifies how >> as the fact that it is in the control mode page is already stated in this paragraph and there is no reason the give the long version of the field name as that is done

no where else.

IBM-045

PDF pg 97, pg 79, 5.9.2.2 Handling new tasks from the faulted initiator port

when ACA is in effect, Table 25

Add in a new column under << new task properties >> for the TASK_ONLY bit. If

ACA attribute and set to 0 then process the task; If ACA attribute and set to

1 then terminate the task; If not ACA attribute then don't care and terminate

the task. This will be in a upcoming proposal.

IBM-046

PDF pg 98, pg 80, 5.9.2.3.2 Handling new tasks from non-faulted initiator ports when ACA is in effect, Table 26

The title of column 1 << TST field value in control mode page >> should be <<

TST >> as the location of the TST is established in the paragraph above the table and the suggest change is how you do it in other tables in SAM-3.

IBM-047 Technical

PDF pg 99, pg 81, 5.9.4 Incorrect logical unit selection, Item b) B)

I have no idea what this statement means << or is not operational when the peripheral device is not ready. >> This needs to be fixed.

IBM-048

PDF pg 100, pg 82, 5.9.6 Sense data, 1st paragraph

The statement << applicable command set standard and applicable >> should be

<< applicable command set standard, and applicable >> (i.e., it is missing a comma).

IBM-049 Technical

PDF pg 101, pg 83, 5.9.7 Unit Attention condition, 4th paragraph

The statement << 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. >> is duplicated in SPC-3 and should be removed from one. I suggest removing to from SPC-3 and placing a pointer to SAM-3 in SPC-3.

IBM-050 Technical

PDF pg 101, pg 83, 5.9.7 Unit Attention condition, 5th paragraph

The statement << 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.

>> is duplicated in SPC-3 and should be removed from one. I suggest removing

to from SPC-3 and placing a pointer to SAM-3 in SPC-3.

IBM-051

PDF pg 101, pg 83, 5.9.7 Unit Attention condition, 5th paragraph

The description of the REQUEST SENSE command handling of Unit attention is not

duplicated in SPC-3 but there is no reference to SAM-3 in SPC-3. This should

be fixed in SPC-3 unless it is decided to move all this into SPC-3. (I know this isn't really a SAM-3 comment but the two are closely linked in this case.

IBM-052

PDF pg 101, pg 83, 5.9.7 Unit Attention condition, Last paragraph

The statement << If the UA_INTLCK_CTRL field in the Control mode page contains 10b or 11b, the logical unit >> should be << If the UA_INTLCK_CTRL field contains 10b or 11b, the logical unit >> as the fact that it is in the

control mode page is already stated in this paragraph

IBM-053 Technical

PDF pg 103, pg 85, 6.2 Establishing a unit attention condition subsequent to

detection of an event, 1st paragraph

The statement <<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. >> is unclear at best and serves no purpose. It should be deleted

along with the last column in table 27 for the same reason.

IBM-054 Technical

PDF pg 104, pg 86, 6.2 Establishing a unit attention condition subsequent to

detection of an event, Last paragraph

The statement << 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. >> should be << Otherwise, the logical unit shall

use one of the other additional sense codes (e.g., POWER ON OCCURRED) when establishing a unit attention condition. >> for the same reason as stated in previous comment.

IBM-055 Technical

PDF pg 104, pg 86, 6.3.2 Hard reset, item b

In the statement <<A part of the response to a power on condition (see 6.3.1).

>> it is not at all clear what a part of the response is a part of. This needs

to be fixed or deleted.

IBM-056 Technical

PDF pg 105, pg 87, 6.3.3 Logical unit reset, Item b

In the statement << A part of the response to a hard reset condition (see 6.3.2). >> it is not at all clear what a part of the response is a part of. This needs to be fixed or deleted.

IBM-057 Technical

PDF pg 107, pg 89, 7.1 Introduction

In this subclause the following statement needs to be added << All supported

task management functions received by a logical unit shall be processed even

if an ACA condition has been established. >>. Note that this means an CLEAR ACA can be issued by any initiator.

IBM-058

PDF pg 115, pg 97, 8.4 Task management events, 2nd state description

The statement << If the TST field in the Control mode page equals 000b, >> should be << If the TST field equals 000b, >> once a subclause is more than enough.

IBM-059

PDF pg 115, pg 97, 8.4 Task management events, 2nd state description

The statement << If the TST field in the Control mode page equals 001b, >> should be << If the TST field equals 001b, >> once a subclause is more than enough.

IBM-060 Technical

PDF pg 118, pg 100, 8.6.2 Simple task, 1st paragraph

The statement << the enabled task state until all older head of queue tasks and older >> is not correct as a head of queue task that comes in after the simple task may move ahead of it in the queue. It should be stated as << the enabled task state until all head of queue tasks and older >>

IBM-061 Technical

PDF pg 118, pg 100, 8.6.3 Ordered task

The statement << shall not enter the enabled task state until all older tasks in the task set have ended >> is not correct as it does not say anything about head or queue tasks that could move ahead of the ordered task. It should be stated as << shall not enter the enabled task state until all older tasks and any head of queue tasks in the task set have ended >>

IBM-062 Technical

PDF pg 120, pg 102, 8.8 Task state transitions, Figure 39

The statement << For simple tasks, all older head of queue and all older ordered tasks >> should be << For simple tasks, all head of queue and all older ordered tasks >> for the reason stated in previous comment.

IBM-063

PDF pg 120, pg 102, 8.8 Task state transitions, Figure 39

The statement <<For ordered tasks, all older ordered tasks have ended. >> should be << For ordered tasks, all older ordered tasks and any head of queue tasks have ended. >> for the reasons stated in previous comment.

IBM-064 Technical

PDF pg 121, pg 103, 8.8 Task state transitions, transition S1:S2 item a

The statement << task state when all older head of queue and older >> should

be << < task state when all head of queue and older >> of the reason stated in

previous comment.

IBM-065 Technical

PDF pg 121, pg 103, 8.8 Task state transitions, transition S1:S2 item b

The statement << state when all older tasks (see 8.4) have ended. >> should be

<< state when all older tasks and any head of queue tasks (see 8.4) have ended. >> of the reason stated in previous comment.

IBM-066

PDF pg 121, pg 103, 8.8 Task state transitions, transition S1:S2 item a

The statement << If the TST field in the Control mode page contains 001b, then

dormant >> should be << If the TST field contains 001b, then dormant >> as one time is more than enough.

IBM-067 Technical

PDF pg 124, pg 106, 8.9.3 Ordered tasks, table 32

The statement << until all older tasks have ended. >> should be << until all

older tasks and head of queue task have ended. >> for the reason stated in previous comment.

IBM-068 Technical

PDF pg 124, pg 106, 8.9.3 Ordered tasks, table 32

The statement << all older head of queue and older ordered tasks have ended.

>> should be << all head of queue and older ordered tasks have ended. >> for

the reason stated in previous comment.

Comments attached to Yes ballot from Mark Evans of Maxtor Corp.:

Maxtor #1
PDF Page 10
Change, "...the term name and world wide identifier..." to, "...the terms name and world wide identifier..."

Maxtor #2
PDF Page 14
Delete the word "usually" in the following sentence: "Well known logical units allow an application client to issue requests to receive and manage specific information usually relating to a SCSI target device."

Maxtor #3
PDF Page 16
I think it would be beneficial if the description of numeric conventions became more consistent from T10 standard to T10 standard. I recommend a combination of what is in the latest draft of SAS-1.1 and SPC-3.

Maxtor #4
PDF Page 17
Change, "[input-1] [,input-2]" to, "[input-1], [input-2]".

Maxtor #5
PDF Page 18
Delete the word "logically" in item (b) in the bulleted list.

Maxtor #6
PDF Page 19
Change, "...the overall SCSI Architecture model;" to, "...the overall SCSI architecture model;"

Maxtor #7
PDF Page 20
Change, "...which includes transmission of the request and response to/from the remote server." to, "...which includes transmission of the request to and the response from the remote server."

Maxtor #8
PDF Page 30
Change, "SCSI port identifier..." to, "The SCSI port identifier..."

Maxtor #9
PDF Page 30
Change, "The SCSI port identifier is equivalent to SCSI identifier." To,

"The SCSI port identifier is equivalent to the SCSI identifier."

Maxtor #10

PDF Page 38

The term, "...peripheral device..." is introduced for the first time in this sentence without any definition. This term is used farther down this page and extensively on the page numbered 81 (pdf 99 of 132). I recommend that an "e.g." be added here or an entry be made in the definitions describing this term.

Maxtor #11

PDF Page 38

Change, "...entities that need LUNs..." to something like, "...entities that are assigned LUNs..."

Maxtor #12

PDF Page 61

As this bit was last described in SAM, I recommend that the following sentence be deleted "Bit 1 provides an obsolete way to request interrupts between linked commands."

Maxtor #13

PDF Page 62

This is the first of many instances of the phrase "...an unit attention..." in this document. Change this phrase in all cases to the phrase "...a unit attention..."

Maxtor #14

PDF Page 64

Change list item 4)A): "for any reason not listed in 1);" to, "for any reason not listed in 2);"

Maxtor #15

PDF Page 68

I recommend that, "A SCSI initiator device shall support a resolution of one byte." be changed to something like, "A SCSI initiator device shall support a resolution of a minimum of at least one byte."

Maxtor #16

PDF Page 71

Change, "...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." to, "...the lifetime of a task or command as it appears to the application client is different from the lifetime observed by the device server."

Maxtor #17

PDF Page 76

Change, "...the application client may request that device server alter command processing..." to, "...the application client may request that the device server alter command processing..."

Maxtor #18

PDF Page 77

Change, "...the application client may request that device server alter command processing..." to, "...the application client may request that the device server alter command processing..."

Maxtor #19
PDF Page 79

Change, "The processing of specific commands (e.g., PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action) the from SCSI initiator ports..." to, "The processing of specific commands (e.g., PERSISTENT RESERVE OUT command with a PREEMPT AND ABORT service action) from the SCSI initiator ports..."

Maxtor #20
PDF Page 81

Change, "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." to, "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 the overlapped command."

Maxtor #21
PDF Page 82

Change, "...and applicable SCSI transport protocol standard." to, "...and the applicable SCSI transport protocol standard."

Maxtor #22
PDF Page 89

Change, "The task management function names are summarized in table 28." to, "The task management function names are listed in table 28."

Maxtor #23
PDF Page 98

Change, "The task management function names are summarized in table 29." to, "The task management function names are listed in table 29."

Maxtor #24
PDF Page 101

Change, "A difference in task priority between tasks does not necessarily override..." to "A difference in task priority between tasks may not override..."

Maxtor #25
PDF Page 104

Change, "In snapshot 1 the task set initially contains one head of queue and one simple task." to, "In snapshot 1 the task set contains one head of queue and one simple task."

Maxtor #26
PDF Page 106

I'm not sure why task 5 is included as a cause for the boundary. I think that this should be deleted.

Maxtor #27
PDF Page 106

I'm not sure why task 5 is included as a cause for the boundary. I think that this should be deleted.

Maxtor #28
PDF Page 106

I'm not sure why task 5 is included as a cause for the boundary. I think that this should be deleted.

Maxtor #29

PDF Page 107

I'm not sure why task 3 is included as a cause for the boundary. I think that this should be deleted.

Maxtor #30

PDF Page 107

I'm not sure why task 3 is included as a cause for the boundary. I think that this should be deleted.

Maxtor #31

PDF Page 107

I'm not sure why task 3 is included as a cause for the boundary. I think that this should be deleted.

Comments attached to Yes ballot from Charles Binford of Sun Microsystems, Inc.:

SAM-3 r17 Comments

Charles Binford

Sun Microsystems, Inc.

Sun 1

(same text in two places)

4.13.3, Page 45, last sentence on page.

4.13.5, Page 47, Last sentence of paragraph following figure 20

Current text:

'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).'

The availability of the same LU through multiple ports is discovered by matching LU names, not SCSI port names.

Sun 2

5.3.1, Page 63, 'ACA ACTIVE' description

The a), b), c), list is overly simplified and not consistent with the tables

in sections 5.9.2.2 and 5.9.2.3 (e.g. Under certain conditions a BUSY is supposed to be returned instead of ACA Active). I suggest that instead of trying to fix this text the section be reworded along the following lines:

'The status shall be returned under certain conditions when an ACA exits within the task set. See 5.9.2.2 and 5.9.2.3 for details.'

Sun 3

5.5, Page 71, a) through f) list of conditions that indicate tasks no longer exist

This list should include receiving the Sense data

Aborted-Command/Overlapped-Commands-Attempted. (See 5.9.3)

Sun 4

5.7.1, Page 72, bottom of page

Sending an Overlapped Command is an action that causes the sending initiator

port's IOs to be aborted. However, I'm not sure if it fits the context since

the condition is an indication of a bug. I brought it up, I'll let the editor/committee decide if it fits.

Sun 5

Multiple sections, multiple pages, but 5.7.2, 5.7.3 on page 73 are examples.

The term 'initiator port' is used in many, many places throughout the document

to describe a subset of tasks, e.g. 'task(s) of another SCSI initiator port...'. My contention is that in most of these cases the group of tasks is

more accurately described as the tasks from a SCSI initiator port sent across

a particular I_T nexus. I already have a document addressing this concern, 04-088, and will continue to pursue this issue via that document instead of in

this set of comments.

Sun 6

5.9.4, Page 81, bottom of page

'Any command except Request Sense or Inquiry:' is incomplete. The command REPORT LUNS needs to be added to this section.

Sun 7

5.9.7, Page 82, sentence e)

'e) Tasks for this SCSI initiator port were cleared by another SCSI initiator

port;'

This is true only if the TAS bit in the control mode page is 0.

Sun 8

5.9.7, page 83, multiple paragraphs

'an unit attention' should be 'a unit attention'.

Comments attached to Yes ballot from Paul D. Aloisi of Texas Instruments:

Editors notes should be resolved before a document goes out for review. Byte limitations should remain the rule in this version of SAM, future versions may want consider this again. There has to be a basic unit definition, the byte should remain the basic unit.

Comments attached to Yes ballot from Roger Cummings of Veritas Software:

VERITAS 1

PDF pg 32, pg 14, 3.2 Acronyms

Add HiSUP to acronyms

Proposed Resolution:

HISUP Hierarchy Supported (see 4.9.1 and SPC-3)

VERITAS 2

PDF pg 39, pg 21, 4.3 The SCSI client-server model

a single command, series of linked commands, or task management function

Proposed Resolution:

a single command, a series of linked commands, or a task management function

VERITAS 3

PDF pg 42, pg 24, 4.5 SCSI Domain

A SCSI device containing logical units that process commands is called a SCSI target device.

Proposed Resolution:

A SCSI device containing one or more logical units that process commands is called a SCSI target device.

VERITAS 4

PDF pg 42, pg 24, 4.5 SCSI Domain

providing a mechanism through which application clients and device servers communicate (see 4.6).

Proposed Resolution:

providing a mechanism through which application clients and device servers or task managers communicate (see 4.6).

VERITAS 5

PDF pg 48, pg 30, 4.7.3 SCSI target/initiator device

"A logical unit is the object to which commands are sent." In 4.7.2 the term "addressed" is used instead of sent, should be consistent

Proposed Resolution:

A logical unit is the object to which commands are addressed.

VERITAS 6

PDF pg 48, pg 30, 4.7.4 SCSI port identifier

The SCSI port identifier is equivalent to SCSI identifier.

Proposed Resolution:

The SCSI port identifier is equivalent to a SCSI identifier (see 3.1.9.3 and Annex B).

VERITAS 7

PDF pg 48, pg 30, 4.7.5 relative port identifier

Question: The last paragraph implies relative port identifiers do not change across power cycles - is this true?

Proposed Resolution:

Clarify

VERITAS 8

PDF pg 50, pg 32, 4.8 Logical Units

A logical unit number is a field (see 4.9)

Proposed Resolution:

A logical unit number (see 4.9) is a field

VERITAS 9

PDF pg 50, pg 32, 4.8 Logical Units

If any logical unit within the scope of a SCSI target device includes dependent logical units in

Proposed Resolution:

If any logical unit within the scope of a SCSI target device includes one or more dependent logical units (see 3.1.23) in

VERITAS 10

PDF pg 50, pg 32, 4.8 Logical Units

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

Proposed Resolution:

The device server is the object that processes the operations received in Device Service Requests.

VERITAS 11

PDF pg 51, pg 33, 4.9.2 LUN 0 address

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

Proposed Resolution:

To address the LUN 0 of a SCSI device the peripheral device address method (see 4.9.6) shall be used.

VERITAS 12

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure

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

Proposed Resolution:

Table 1 describes a single level subset of the format described in 4.9.4 for logical unit numbers 255 and below.

VERITAS 13

PDF pg 52, pg 34, 4.9.3 Single level logical unit number structure

Comment: dependent logical units are defined in clause 3, extended addressing logical units are not defined

Proposed Resolution:

Add definition for extended addressing logical units to clause 3

VERITAS 14

PDF pg 55, pg 37, 4.9.6 Peripheral device addressing method

Question: is the term relay here common usage? Should it be more clearly defined?

Proposed Resolution:

Clarify

VERITAS 15

PDF pg 56, pg 38, 4.9.7 Flat space addressing method

Question: Why is the last paragraph on page 38 included for this address

method and not for any others?

Proposed Resolution:

Make consistent

VERITAS 16

PDF pg 57, pg 39, 4.9.8 Extended logical unit addressing

(see table 6 in 4.14)

Proposed Resolution:

(see table 6 in 4.9.4)

VERITAS 17

PDF pg 57, pg 39, 4.9.8 Extended logical unit addressing

Comment: Table 11 is correct but confusing e.g. the 1st line contains a length of one byte but references table 12 which is the "two byte extended addressing format".

Proposed Resolution:

Make the length column in table 11 numeric i.e. make the heading identify bytes as the value and put the 1st line entry as "1" etc.

VERITAS 18

PDF pg 58, pg 40, Table 15 Eight byte extended logical unit addressing format

Question: Why is the byte numbering different in this table from the three preceding ones?

Proposed Resolution:

Number the three lines, n, n+1 and n+6

VERITAS 19

PDF pg 59, pg 41, 4.9.9 Well known logical unit addressing

A SCSI target device may support zero or more well known logical units.

Proposed Resolution:

A SCSI target device may support zero or more well known logical units
(see 4.10).

VERITAS 20

PDF pg 60, pg 42, 4.10 Well known logical units

Well known logical units are addressed using the well known logical unit addressing method of extended logical unit addressing (see 4.9.8).

Proposed Resolution:

Well known logical units are addressed using the well known logical unit addressing method (see 4.9.9).

VERITAS 21

PDF pg 60, pg 42, 4.10 Well known logical units

shall follow the rules for selection of invalid logical units described in 5.9.4.

Proposed Resolution:

shall follow the rules for selection of incorrect logical units described in 5.9.4.

VERITAS 22

PDF pg 60, pg 42, 4.10 Well known logical units

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

Proposed Resolution:

A SCSI target device may have more than one SCSI target device name if the SCSI target device supports multiple SCSI transport protocols (see 4.7.2 for additional requirements related to name formats).

VERITAS 23

PDF pg 68, pg 50, Note 4

exclusive access reservation

Proposed Resolution:

exclusive access persistent reservation

VERITAS 24

PDF pg 68, pg 50, Note 4

Question" SPC-2 is included in the Normative References. Should reserve/release also be addressed in SAM-3, in Note 4 and in RESERVATION CONFLICT in 5.3.1 (PDF pg 81, pg 63)??

Proposed Resolution:

Clarify

VERITAS 25

PDF pg 69, pg 51, 4.14 Model for dependent logical units

When the dependent logical unit model is utilized, the hierarchical logical unit structure defined in 4.9.4

Proposed Resolution:

When the dependent logical unit model is utilized, the hierarchical logical unit number structure defined in 4.9.4

VERITAS 26

PDF pg 70, pg 52, 4.14 Model for dependent logical units

Question: Does this a) thru d) lists here contradict with the requirement on the previous page to use the hierarchical logical unit structure defined in 4.9.4?

Proposed Resolution:

Clarify

VERITAS 27

PDF pg 80, pg 62, 5.2 Command descriptor block (CDB)

Comment: SBC is not included in the Normative References

Proposed Resolution:

Add SBC to Normative References

VERITAS 28

PDF pg 28, pg 10, 3.1.72 port

Add reference to Annex B

Proposed Resolution:

(see 3.1.96 and Annex B)

***** End of Ballot Report *****