

Date: May 22, 1995

From: Charles Monia
SAM Technical Editor
Digital Equipment Corporation
334 South Street
Shrewsbury, MA 01752

Voice: (508) 841-6757

EMAIL: monia@shr.dec.com

To: Members of X3T10

Subject: X3T10/95-202R2 Responses to SAM Public Review Comments

The following are responses to the public review comments received from: George Penokie of IBM, Lansing J. Sloan of LLNL, Tom Wicklund of Intellistor and Rodney Van Meter of the USC Information Sciences Institute.

The resolution of these comments as reflected in this document was discussed during the SCSI Working Groups of March and May, 1995 and approved by X3T10 at the 11 May 1995 Plenary.

In this revision of this document (revision 2), the following responses have been modified from revision 1:

LLNL #008
LLNL #012
LLNL #016
LLNL #017
LLNL #018
LLNL #059

Each response has the format:

#nnn (?) y.y.y (clause title), Description

Where:

nnn is the comment number assigned by the originator (if applicable);
? is the type (E: Editorial, T: Technical);
y.y.y is the referenced section number

Proposed Responses to LLNL Comments

#001 (E) Comment on 4.1.101 (unlinked command definition)

Comment:

It appears that the last command in a series of linked commands is defined to be an unlinked command. Is that the intent?

Response:

That is correct. However, no change to the definition is warranted.

#002 (E) Comment on 4.3 Acronyms and Abbreviations

Should entries for "SCSI-1" and "SCSI-3" be added, or "SCSI-2" deleted?

Response:

The manner in which SCSI references are cited and referred to must be consistent across all SCSI-3 documents. The technical editors will define a common set of conventions for such references.

#003 (E) Comment on 5.2.1, last paragraph (The SCSI Client-Server Model)

Most of this paragraph should be in Clause 5.3 rather than in 5.2.1.

Response:

Comment accepted.

#004 (E) Comment on 5.2.1 (The SCSI Client-Server Model), last paragraph, last sentence

The second word should be "ensure," not "insure". A global search for this wording throughout the SAM document may be in order.

Response:

Comment Accepted.

#005 (T) Comment on 5.2.1 (The SCSI Client-Server Model), last paragraph

Somewhere, SAM should specify information such as the following. "An SCSI-3 protocol standard shall specify formats and encodings of addresses and identifiers (e.g., in the request's, command descriptor blocks, parameter blocks, etc.)." It's not clear that this is a complete list of such items that protocol standards must specify. Perhaps this information could be added after the fourth sentence of the last paragraph of 5.2.1.

Response:

This matter should be addressed in the next version of the standard (SAM -2).

#006 (E) Comment on 5.3, Figure 8

This figure shows the box for "Service Delivery Port" having two parents. Should Clause 4.6.4 mention the possibility of multiple parents in such diagrams?

Response:

Comment accepted. The text will be modified to reflect the cited figure.

#007 (E) Comment on 5.5.1 (Synchronizing Client and Server States)

At least four terms, "client," "initiator," "server," and "target," are used in this clause. Is it reasonable to replace all instances of "initiator" by "client" and all instances of "target" by "server"? Or other consistent replacements?

If "client" and "server" are the best terms, then "initiator" at the end of sentence 2 in paragraph 1 probably should be replaced by "target." By contrast, the use of "target" in sentence 3 seems to be reasonable. The last sentence of that paragraph uses both "target" and "initiator," possibly correctly. The last sentence of the second paragraph should say "client" rather than "initiator."

Response:

Comment rejected. The terms "client" and "server" are used when referring to a generic functional relationship between distributed objects which is true independently of SCSI. The terms initiator and target refer to that relationship in the context of SCSI. The usage in the cited paragraph is consistent with that distinction.

#008 (T) Comment on 5.5 (The Service Delivery Subsystem), last paragraph, first sentence

What is meant by "The service delivery subsystem provides error-free transmission of requests and responses ..."? It does not seem likely that SAM prohibits errors or requires protocol standards to prohibit errors. (Possibly, some text similar to the last paragraph of 5.5.2 would be appropriate, to explain that SAM assumes error-free transmission for simplicity but, say, requires protocol standards to provide a certain level of assurance of error detection and recovery. What are SAM's actual requirements?)

Response:

To eliminate the implied requirement, the cited paragraph will be reworded as follows:

"The service delivery subsystem is assumed to provide error free transmission of requests and responses.

#009 (T) Comment on 6.1 (Command Descriptor Block), first paragraph, last sentence

The last sentence is true, but incomplete. Consider adding a sentence such as the following: "Protocol-specific formatting of parameters is specified in the applicable SCSI-3 protocol standards."

If the last sentence is normative (i.e., a requirement for command standards), then it should be reworded, perhaps by changing "... are specified ..." to "... shall be specified ...". The new sentence proposed in the preceding paragraph should then be similarly modified.

(FCP Rev 10, clause 5.2, is an example of a protocol standard specifying protocol-specific formatting of a parameter.)

Response:

Comment rejected. The applicable SCSI-3 command standard is the appropriate place to address the issue of protocol-specific arguments.

#010 (T) Comment on 6.1.2 (Control Field), link bit

SCSI-2 has text in clause 7.8.2 that concerns the interaction of queuing and linked commands. "A series of linked commands constitute a single I/O process. .. A command received with a HEAD OF QUEUE TAG message shall not suspend a series of linked commands for which the target has begun execution."

Is that, or a similar, requirement supposed to be in SCSI-3? If so, is SAM the appropriate standard that should state the requirement? (I have not found such text in any SCSI-3 document, though could easily have failed to see it.)

Response:

Comment rejected.

This removal of this language from the SCSI-3 queuing model was intentional. The goal of the new model is to specify queuing only in terms of externally observable behavior. The SCSI-2 example cited in the comment describes behavior that is internal to the target and therefore not appropriate for inclusion in SAM.

#011 (T) Comment on 6.1.2 (Control Field), link bit

Is it supposed to be possible in SCSI-3 to use linked commands to create complex atomic operations? Assume there may be multiple initiators. As a specific example, suppose one initiator has reserved two adjacent extents of a logical unit. Suppose that initiator links three commands. The first two commands each release a reservation. The third reserves the combined extent. For this example, assuming no errors occur, is there any assurance that no commands are accepted that access the extents while the extents are unreserved? Alternatively, are there specific sequences that permit the extents to be accessed before the third (the RESERVE) command is executed?

Text (perhaps informative) discussing such issues seems useful.

Response:

Comment rejected. The use of linked commands in the manner described above is not supported by the architecture. Adding an explicit disclaimer now would require a fairly precise explanation of "atomicity", with the strong possibility of sowing further confusion and adding months to the review and approval cycle. For that reason, the change should be deferred to SAM -2. In the meantime, a technical information bulletin can be drafted if necessary.

#012 (T) Comment on 6.1.2 (Control Field), vendor-specific

There should be some description of "vendor-specific" bits. For instance, should SAM say the value "00" means there is no vendor-specific information?

Response:

Comment accepted.

The term will be added to the glossary.

#013 (E) Comment on 6.2 (Status), first paragraph, second sentence

This sentence mentions a service response of "Command Complete," consistent with the service response "Send Command Complete" in clause 6.3 but not consistent with "Task Complete", the first of the listed "Service Response" values in clause 6. We suspect clause 6.2 should say "Task Complete."

Response:

Comment accepted. The service response should be Task Complete.

#014 (T) Comment on 6.2 (Status), paragraph on "BUSY"

Note: this comment sneaks in some issues related to Livermore's desire for "distributed I/O" in large configurations where initiators are not all equally trusted. The paragraph describes what status to return if a command is received from an "otherwise acceptable" initiator but says nothing about what status to return if a command is received from an unacceptable initiator.

Response:

Comment rejected. The text states that an acceptable initiator is one having no reservation conflict. The standard is clear on the status to be returned when such a conflict exists.

#015 (E) Comment on 6.2 (Status), paragraph on "RESERVATION CONFLICT"

The first sentence mentions a "RESERVE UNIT" command. Since that command is no longer in SPC, but new commands are being proposed, please reword the paragraph to refer to all reservation commands. Consider making the following changes inside the parentheses: change "see" to "such as", delete "and RESERVE UNIT", and change "commands" to "command".

Response:

Comment accepted.

#016 (T) Comment on 6.3.1 (Data Transfer Protocol Service), third-from-last paragraph

The first sentence of this paragraph probably is violated by most "initiator implementations" inside peripherals that implement the COPY command, since when they behave as initiators they do not support a resolution of one byte. It is not obvious how to improve the wording while keeping the requirement as strong as seems intended. Perhaps add: "Coarser resolutions are permitted during the execution of some commands in command standards (e.g., the COPY command in SPC)."

Response:

Comment rejected.

The draft standard reflects the intention of the committee.

#017 (T) Comment on 7.1 (Abort Task), last paragraph, last sentence

If a service delivery subsystem misorders (as is specifically allowed by the last paragraph of clause 5.5.2), the target can guarantee that no further responses are sent from the task, but the target cannot guarantee that the client will receive no responses from the task after that client receives the response to ABORT TASK. The current text appears to be correct (depending on what "further" means) but misleading, and therefore should be reworded to point out that clients can receive task responses after they receive the service response to ABORT TASK.

Response:

Comment accepted.

The last sentence will be reworded as follows:

"In either case, after sending the Task Command Complete response, the target shall guarantee that no more responses from the aborted task are sent."

#018 (T) General comment

In large SCSI configurations where all initiators are not equally trusted, peripherals should be allowed to reject in some manner unwanted commands from untrusted initiators. Such issues are presently beyond the scope of SCSI standards. However, SAM should not have words that conflict with such possible future configurations. In a quick scan of SAM Revision 016, no offending text was detected. If there are any changes to SAM, they should not introduce offending text, and if any now exists, it should be corrected.

Response:

Comment rejected.

The comment author should raise these issues as specific change proposals to be considered by the committee SAM -2.

#019 (E) Comment on Table of Contents

Each of clauses 2.1, 4.4.1, 5.2.1, 5.6.2.1, and 8.2.1 is not followed by another clause at the same level. In "Object definition (6)", the letter "d" should be upper case. "Figure 23" should be followed by a colon. Any corrections made here apply also to the body of the document.

Response:

Comment accepted.

#020 (E) Comment on 4.1.48 (definition of "option, optional")

It seems desirable to append to the second sentence "if compliance is claimed", or to find some other change for the definition of "option."

Response:

Comment rejected. The text in SAM is identical to the accepted definition for this term that is used by all SCSI-3 standards.

#021 (E) Comment on 4.1.51 (definition of "pending task")

Change the "pending task" definition by adding the text ".. nor a completed task."

Response:

Comment rejected. Task completion is an event, not a task state. Tasks are deleted upon completion.

#022 (E) Comment on 4.1.55 (definition of "protocol-specific")

The definition of "protocol-specific" talks about a "SCSI-3 protocol standard," but the term "protocol standard" is never defined. Clauses 2 and 2.1 do talk about "implementation standards" and could easily introduce terms such as "command standard," "protocol standard," and "interconnect standard," if such terms would be helpful. Clause 6.6.4 mentions "device command standard" and "protocol standard."

Response:

Comment accepted. A cross reference will be added to the list of SCSI-3 protocol standards.

#023 (E) Comment on 4.1.55 (definition of "protocol-specific")

The pointer to clause 1 seems bad. Clause 2 or 2.1 seems better.

Response:

Comment accepted. The cross reference will be corrected.

#024 (E) Comment on 4.1.55 and 4.1.57 (definition of "protocol-specific" and "protocol option")

The first of these uses "protocol standard," the second uses "protocol specification." Perhaps the latter is preferable unless "protocol standard" is defined.

Response:

Comment accepted. All such references will be changed to "SCSI-3 protocol standard" and a cross reference will be added to the appropriate clause where such standards are listed.

#025 (E) Comment on 4.1.57 (definition of "protocol option")

Why is "protocol option" defined but neither "interconnect option" nor "command option"?

Response:

Neither term is used in SAM.

#026 (E) Comment on 4.1.96 (definition of "task set")

Spelling: Use "dependent." Probably a spell checker should be run on the document and questionable spellings brought to the committee.

Response:

Comment accepted.

#027 (E) Comment on 4.3 (Acronyms and Abbreviations)

For "SCSI," replace "Either" with "Any one of".

Response:

Comment accepted.

#028 (E) Comment on 4.4 (Editorial Conventions), second paragraph, first sentence

The word "attributes" should be singular.

Response:

Comment accepted.

#029 (E) Comment on 4.6 (Objects and Object Notation), first paragraph, third sentence

SAM does not define an "I/O bus" object; the example is poor.

Response:

Comment accepted, "I/O bus" will be changed to "physical interconnect".

#030 (E) Comment on 4.6.1

The definition of symbol "nn" is confusing. An example that includes at least one digit greater than 1 would help. Currently the confusion is not resolved until clause 4.6.2.

Response:

Comment accepted.

#031 (E) Comment on 5.1 (Introduction), second normal paragraph, last sentence

It is not clear whether that sentence ("The description of internal behavior ...") applies to SAM or whether it applies to all SCSI-3 standards. It almost surely should not apply to CAM.

If it is intended to apply to any SCSI-3 standards other than SAM, that intent should be stated more clearly.

Response:

Comment accepted. The sentence will be reworded to limit applicability to SAM.

#032 (E) Comment on 5.1 (Introduction), last paragraph

This paragraph is helpful. However, a complete list of SAM clauses that apply to any SCSI-3 standards would be useful. Clauses 4.5, 5.2.1 (at least the last paragraph), 5.6.2.1 (on addressing the task manager), 6.1, and 6.2 should be included in such a list, in addition to 6.3 and 7.7.

Response:

Comment rejected. A complete listing of such clauses is impractical and would be too long to be of use.

#033 (E) Comment on 5.2, first paragraph, second sentence

"Dashed horizontal arrows" here conflicts with "dotted" in the paragraph following figure 18.

Response:

Comment accepted.

#034 (E) Comment on Figures 5, 6, 16(?), 18, and maybe others

The quality of figures is low enough that it is hard to tell whether lines are dashed, dotted, solid, or something else.

Response:

Comment accepted.

The line drawings will be reworked.

#035 (E) Comment on 5.3, 5.4 (The SCSI Structural Model, SCSI Domain) and maybe elsewhere

The first paragraph in 5.3 seems indented by a blank. The description for "Service Delivery Subsystem" in 5.4 seems indented by a blank. The document should be scanned for other such possibly unwanted blanks.

Response:

Comment accepted.

#036 (E) Comment on 5.4 (SCSI Domain)

Is "Service Delivery Interface" in Figure 9 (or any other places that it might occur in SAM) the same as "Service Delivery Port"?

Response:

References to "service delivery interface" are incorrect and will be replaced by "service delivery port".

#037 (E) Comment on 5.5.1 (Synchronizing Client and Server States), first paragraph, last sentence

Insert an apostrophe before the "s" in "initiators".

Response:

Comment accepted.

#038 (E) Comment on 5.5.2, fourth paragraph, last sentence

"or places" should be "nor places", and after "requirement on" there should be a comma.

Response:

Comment accepted.

#039 (E) Comment on 5.6, under Object Definition (3)

The text describing a Target should end with a period (or else a lot of periods should be removed from similar descriptions in SAM).

Response:

Comment accepted.

#040 (E) Comment on 5.6 (SCSI Device Models)

For the description of "Service Delivery Port," the words "interconnect subsystem" probably should be "service delivery subsystem."

Response:

Comment rejected. The intended reference is to the interconnect subsystem

#041 (E) Comment on 5.6 (SCSI Device Models)

In the text describing "Port Identifier", the words "by the device" probably should be deleted or explained. For instance, in Fibre Channel it is quite possible that the identifier is a Fibre Channel address assigned by a Fabric and that the device has no control at all.

Response:

Comment accepted. The text will be reworded to read:

"Port Identifier: The identifier assigned to the port."

#042 (E) Comment on 5.6.2 (SCSI Target)

In the description of "Logical Unit 0", if there should be two sentences then "see" should be "See" but if there should be one sentence then the period after "zero" should be removed. In either case, the period after "5.6.3" and before the closing parenthesis should be deleted.

Response:

Comment accepted.

#043 (E) Comment on 5.6.2.1 (The Task Manager), first paragraph, second sentence

Text should be added to explain which, if any, LUN value is or can be used when external ports communicate with the task manager. Possibly clause 7 is a better place to provide the information.

Response:

Comment accepted. However, as stated in the text, the address of the task manager is the target identifier, which does not contain the logical unit number. i.e.. There is one task manager per target. cross references to the hierarchy diagram of figure 14 and the object definitions for the target and target identifier will be added.

#044 (E) Comment on 5.6.2.1 (The Task Manager), last paragraph

To improve clarity, change "An" to "If an", change "wishing" to "wishes", change "insure" to "ensure", and insert "it" after the comma.

Response:

Comment accepted.

#045 (E) Comment on 5.6.3 (Logical Unit)

The use of the term "Implementation-specific information" and the description of this term are not helpful. Is such information supposed to be standard, or outside the scope of standards? (Keep in mind nearly all SCSI-3 standards are "implementation standards" in clause 2.1.)

Response:

Comment Accepted

The reference to "implementation specific information" will be deleted.

#046 (E) Comment on 5.7 (The SCSI Model For Distributed Communications), first paragraph that follows figure 16

For clarity, insert ", respectively" before the period ending the second sentence. To make the result correct, either interchange "LLP" and "ULP" in the first sentence or else interchange "outgoing" and "incoming" in the second sentence.

Response:

Comment accepted.

#047 (E) Comment on Figure 18

Parts of the right side of the figure seem to be missing.

Response:

Comment accepted. The figure will be corrected.

#048 (E) Comment on 6.1 (Command Descriptor Block), third paragraph

Should "medium" be replaced by "media information," which is in the glossary?

Response:

No, the existing wording is what was intended. The term "media" is used in accordance with its normal English meaning.

#049 (E) Comment on 6.1 (Command Descriptor Block), Table 1

Should there be an alteration to the table to indicate an arbitrary number of bytes between byte 1 and byte "n-1"?

Response:

No, the stylistic conventions for X3T10 standards do not allow the inclusion of ellipses in tables.

#050 (E) Comment on 6.1.2 (Control Field)

The fourth paragraph after Table 3 should be moved to immediately follow Table 3.

Response:

Comment accepted.

#051 (E) Comment on 6.2 (Status), "GOOD"

Should "the command" be "an unlinked command"? Or, perhaps, "an unlinked command or the last of a series of linked commands"?

Response:

The sentence will be reworded as follows:

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

#052 (E) Comment on 6.2, "INTERMEDIATE"

The wording is awkward, in that it says a "successfully completed command" can end with "CHECK CONDITION" or other such status. It should be rewritten. The text in the following paragraph is a better model.

Response:

Comment rejected. The existing wording is identical with the SCSI -2 version. To avoid confusion, such wording, however awkward it might seem, is best retained in the document whenever possible.

#053 (E) Comment on 6.2 (status), "TASK SET FULL"

Should "TASK SET FULL" be prohibited for a CDB in a task other than the first one? Consider linked CDBs.

Response:

Comment rejected. The task is created and entered into the task set upon receipt of the first command. Therefore, the TASK SET FULL status cannot be returned for any except the first in a sequence of linked commands. It is inappropriate for a standard to explicitly prohibit behavior that is logically precluded.

#054 (E) Comment on 6.3 (Protocol Services in Support of Execute Command)

Should the "Send SCSI Command" request have a Command Byte Count, in analogy with DATA IN and DATA OUT requests?

Response:

No, the byte count is included in the CDB.

#055 (E) Comment on 6.3.1 (Data Transfer Protocol Services), last paragraph

Should the last word, "undefined," be "unspecified"?

Response:

Yes, the text will be changed.

#056 (E) Comment on 6.4 (Task and Command Lifetimes), paragraph before list of responses (page 51)

Delete the second colon at the end of the paragraph.

Response:

Comment accepted.

#057 (E) Comment on 6.4 (Task and Command Lifetimes), list of responses

Under b), a comma (not period) should follow "POWER ON".

Response:

Comment accepted.

#058 (E) Comment on 6.4 (Task and Command Lifetimes), list of responses

Under f), does SAM want to discuss "OTHER PORT"?

Response:

No, the reference to "OTHER PORT" will be removed.

#059 (E) Comment on 6.4 (Task and Command Lifetimes)

Should there be a discussion (perhaps elsewhere) of possible hazards if a CDB (not the first of a task) is sent more or less concurrently with a unit attention condition or something else indicating the task was blown away?

Response:

In a full duplex, non-interlocked bus, there are two conditions that need to be considered. In the first, the task sends LINKED COMMAND COMPLETE. In response to this status, the initiator sends the next command in the sequence. While the new command is in transit, an exception causes the task to be aborted. This case always results in a Unit Attention condition. In the committee's opinion, this is covered adequately in SAM.

In the second, an auto contingent allegiance condition occurs after the task sends LINKED COMMAND COMPLETE. Although the task is not aborted, it is unable to complete. When received, the new command becomes 'suspended information' that is not available to the task until the condition is cleared.

This case should be described in SAM -2.

#060 (E) Comment on 6.5.2 (Linked Command Example), paragraph numbered 2

One of the periods after the last sentence should be deleted.

Response:

Comment accepted.

#061 (E) Comment on 6.6.1.2 (Clearing an Auto Contingent Allegiance Condition), fourth paragraph

Delete the second comma.

Response:

Comment accepted.

#062 (E) Comment on 6.6.2 (Overlapped Commands), second paragraph

Why is a tagged task "with a tag value exceeding FFh" treated differently than tagged tasks with smaller tag values? Without some explanation, this looks like a typographical error.

Response:

The value of a duplicate tag is reported in the qualifier associated with an additional sense code of TAGGED OVERLAPPED COMMANDS. Since this is a one-byte qualifier, the occurrence of a duplicate tag exceeding FFh is reported with an additional sense code of OVERLAPPED COMMANDS ATTEMPTED.

Response:

Reword the sentence as follows:

Current wording:

"Otherwise, an additional sense code of TAGGED OVERLAPPED COMMANDS shall be returned with the additional sense code qualifier set to the value of the duplicate tag."

New wording:

"Otherwise, an additional sense code of TAGGED OVERLAPPED COMMANDS shall be returned with the additional sense code qualifier byte set to the value of the duplicate tag."

#063 (E) Comment on 6.6.3 (Incorrect Logical Unit Selection), a), first paragraph, third sentence

The sense data must be returned by "the target", not the nonexistent logical unit.

Response:

Comment accepted.

#064 (E) Comment on 6.6.3 (Incorrect Logical Unit Selection), d), second paragraph

Why is auto contingent allegiance mentioned in case d) but not in any of cases a), b), or c)?

Response:

In cases a), b) the sense data to be returned does not depend on whether or nor an ACA condition is in effect. In case c), the normal rules for sense data apply. What must be defined in case d) is the response to a REQUEST SENSE when no ACA condition is in effect. Consequently, no change to the document is required.

#065 (E) Comment on 6.6.3 (Incorrect Logical Unit Selection), d), second paragraph

The paragraph says what happens unless an ACA exists. What happens if an ACA does exist?

Response:

The normal rules governing the return of sense data apply. No change to the document is required.

#066 (T) Comment on 6.6.3 (Incorrect Logical Unit Selection)

Consider adding a case "e) The target supports the logical unit under miscellaneous, perhaps transient, internal circumstances. For instance, a processor device might use LUNs to identify transactions and might make the LUN exist from the beginning to the end of the transaction and no later."

Response:

Comment rejected.

The above scenario may be handled according to one of the four cases already defined. For example, the rules for case b) or case c) might be appropriate.

#067 (E) Comment on 6.6.4 (Sense Data), first paragraph, last sentence

Before "SCSI-3 protocol standard", insert "applicable."

Response:

Comment accepted.

#068 (E) Comment on 6.6.5 (Unit Attention Condition), fifth paragraph, part 1)

Delete the comma following the last word.

Response:

Comment accepted.

#069 (E) Comment on 6.6.5 (Unit Attention Condition), fifth paragraph, part 2)

Three actions are mentioned. The first (report) is mandatory, the second (discard) is optional, and it's unclear whether the third (clear) is mandatory or optional. Assuming the third is mandatory, it seems better to reorder the three actions to make the two mandatory ones occur first. If two of them are optional, the wording should be changed to be explicit that each of them, independent of the other, is optional.

Response:

Comment accepted. To make the text grammatical, item 2 will be modified as follows:

"2) report the unit attention condition.

If the unit attention condition is reported, the logical unit may discard any pending sense data and may clear the unit attention condition."

#070 (E) Comment on 8 (Task Set Management), third paragraph, second sentence

There should be commas following "TASK SET FULL" and "ACA ACTIVE".

Response:

Comment Accepted.

#071 (E) Comment on 8.1 (Terminology)

Is "task set" ever clearly explained as a concept in SAM? Clause 8.1 would probably be the right place.

Response.

Yes. See the glossary.

#072 (E) Comment on 8.2 (Task Management Events)

In the description of "task abort", should "6.6.2" be "8.2.1" or something else?

Response:

The reference should point to subclause 8.2.1

#073 (E) Comment on 8.3 (Task States)

Should there be an 8.3.5 to discuss "Current"?

Response:

No. "Current" is not a task state.

#074 (E) Comment on 8.6.3 (Task Management Example 3), Figure 27

For the task set in the upper right corner, should there still be an ordered blocking boundary before ORDERED task 3? Should all four parts of the figure show an ordered blocking boundary above ordered task 6? Only the lower right one shows it now.

Response:

Comment Rejected.

The ordered blocking boundary delimits the set of enabled tasks -- that is, the set of tasks that are allowed to complete. The boundary does not apply to tasks whose state prevents them from completing.

#075 (E) Comment on Annex A Tables

It appears that the columns in Table 1, 3, and 4 are not always aligned. It appears that a couple of columns in Table 2 are sometimes misaligned.

Response:

In accordance with the IBM public comment, this annex will be deleted.

#076 (E) Comment on Annex A Tables

More explanation of what the tables are saying would be welcome. For example, in Table 1 does "I1L0H" decode to "initiator 1, LUN 0, Head of queue"? Are events that are successive in time denoted as separate rows as one reads down toward the bottom of the table?

Response:

See comment 075

Proposed Responses to IBM Comments

#001 (T) - Annex A should be removed as it is redundant to the information contained within the body of the standard. If it remains it may cause confusion.

Response:

Comment accepted.

#002 (E) - Page 10 - There are four cases of standards not having their X3T10 number. This should be corrected as all those standards have been assigned numbers.

Response:

Comment accepted.

#003 (E) - Page 22 ; clause 4.7 - The first equation either has one too many or one too few square brackets.

Response:

Comment accepted.

#004 (E) - Page 41 ; Figure 18 - The right side of the figure has been cut off.

Response:

Comment accepted.

#005 (E) -Page 41 ; Figure 18 - The dotted lines are difficult to see as being dotted lines.

Response:

Comment accepted. The line art will be reworked.

#006 (E) -Page 47 ; clause 6.1.2 ; ACA Active paragraph number c

The sentence 'and the ACA bit was set to one' should read 'and the NACA bit was set to one'.

Response:

Comment accepted.

007 (E) - Page 54 ; clause 6.6.1.1 ; third paragraph after notes - The sentence 'If the ACA bit was set to zero' should read 'If the NACA bit was set to zero'.

Response:

Comment accepted.

#008 (E) -Page 55 ; clause 6.6.1.2 ; third paragraph ; This paragraph should read:

"If the NACA bit is set to zero in the CDB control byte of the faulting command, then the SCSI-2 rules for clearing contingent allegiance shall apply. In this case, the logical unit shall also clear the associated contingent allegiance condition upon the return of sense data by means of the autosense mechanism described in clause 6.6.4.2."

Response:

Comment accepted.

#009 (E) -Page 66 ; clause 8 ;third paragraph ; The second sentence 'status of BUSY, RESERVATION CONFLICT, TASK SET FULL ACA ACTIVE or CHECK CONDITION' should be 'status of BUSY, RESERVATION CONFLICT, TASK SET FULL, ACA ACTIVE or CHECK CONDITION'.

Response:

Comment accepted.

#010 (E) - Page 67 ; clause 8.2.1 ; paragraph number d - The control mode page in that sentence should have a cross reference to SPC.

Response:

Comment accepted.

#011 (T) Page 67 - Status priority

The second to last paragraph of section 6.6.5 (revision 16) should be modified as given below since it leads to unnecessary reset operations.

"If an initiator issues a command other than INQUIRY or REQUEST SENSE while a unit attention exists for that initiator (prior to generating the auto contingent allegiance condition for the unit attention condition), the logical unit shall not perform the command and shall report CHECK CONDITION status unless a higher priority status as defined by the logical unit is also pending (see 6.2)."

Subclause 6.2 should be modified as follows:

"If more than one condition applies to a completed task, the report of a BUSY, RESERVATION CONFLICT, ACA ACTIVE or TASK SET FULL status shall take precedence over the return of any other status for that task."

Note:

During the September X3T10 working group, the proposal given above was received favorably. However, to avoid delaying the public review the working group suggested that the change be deferred to SAM -2 or submitted as a public review comment.

Response:

Comment Accepted.

Proposed Responses to Intellistor Comments

#001 (E) Page 42, Section 6, definition of "service response": Input arguments

"Data-Out Buffer" and "Task Attribute" should be swapped in the description at the middle of the page so that the order of the descriptions matches the order in the definition.

Response:

Comment Accepted.

#002 (E) Page 47, Section 6.2, definition of ACA ACTIVE:

Item c) refers to the "ACA bit", it should instead refer to the "NACA bit".

Response:

Comment Accepted.

#003 (E) Page 47, Section 6.3:

In the first paragraph of this section, the second sentence should refer to the "Send SCSI Command Protocol Service request" (add capitalization and the word "request"). This is consistent with the rest of the paragraph's capitalization and use of the words confirmation, indication, and response.

Response:

Comment Accepted.

#004 (T) Page 56, Section 6.6.3: I think items c) and d) should contain identical text in the second paragraph. Item c) specifies that the target should return sense data without specifying the sense data (contrary to the other 3 items in this section).

Response:

Comment rejected.

The two conditions are not identical. In case c), since the unit is known to be non-operational, it may be appropriate to return something other than NO SENSE in response to a REQUEST SENSE command. In case d), the target can't tell whether the unit is broken or not. Also, while c) seems like a special case of item d), retaining the distinction is useful for the sake of clarity.

#005 (T) Page 72, section 8.6.2 (Task Set Management Example 2 -- HEAD OF QUEUE tasks) :

This example (HEAD OF QUEUE tasks) is very confusing: It isn't obvious why a group of simple tasks have an Ordered Blocking Boundary between them (between task 2 and task 4). I think this is because when task 3 is accepted then subsequent simple tasks are dormant until all head or queue tasks complete. It might help the example to show this since anybody who understands why the Ordered Blocking Boundary is present at the start probably doesn't need the example.

Response:

Comment accepted.

Rather than modify the existing example, a new example will be created to show how an HOQ task creates the ordered blocking boundary which partitions the task set.

#006 (T) Page 72, section 8.6.2 (Task Set Management Example 2 -- HEAD OF QUEUE tasks):

The transition from 2 to 3 makes the dormant tasks enabled when Head of Queue task 3 completes. However Head or Queue task 7 is still enabled. It appears from this example that the term "earlier" is used to mean issued earlier in time rather than earlier in the queue (I suggest the term "earlier" be defined here so that it's clear).

Response:

Comment accepted.

#007 (T) Page 72, section 8.6.2 (Task Set Management Example 2 -- HEAD OF QUEUE tasks) :

The definition of Head of Queue seems counter-intuitive here. If a system normally uses only Simple tasks, there is no requirement that Head of Queue tasks be executed before any previously accepted Simple task (since all are in the enabled state). In fact, I don't find anything in the text to imply that a Head of Queue task must be placed at the head of the task queue, merely that it be placed someplace in the queue in enabled state. The requirement that all earlier (in time) Head of Queue tasks be used to determine when a task is dormant means that tasks must be queued both by order of issue and by enabled or dormant. This seems to add a non-obvious burden on implementations.

I realize this has been argued over a lot, so won't suggest a technical change, though renaming "Head of Queue tasks" to "Immediately Enabled tasks" might be more honest. I do suggest that this example be expanded greatly with explanations of why "Head of Queue" doesn't really mean "execute first".

Response

Comment accepted.

In addition to adding another example, the descriptions will be revised to be more clear. The terminology for task attributes, however, will not be changed.

Proposed Responses to USC Information Sciences Institute Comments

The following represents the gist of several private email exchanges.

#001 (T) Clause 8.6.2, Task Management Example 2 -- HEAD OF QUEUE task

| ... was confused by [this example]. As the tasks are numbered in the order received, though, I think even a simple text modification might solve the problem. Change note 1 to read:

1) Head of Queue task 3, received as a Head of Queue task, is in the enabled state. Simple tasks 1 and 2, received and already enabled when Head of Queue task 3 was received, remain enabled. Tasks 4,5 and 6 are dormant, as they were received after Head of Queue task 3, and therefore must wait for it to end.

This is a little long, and does have the disadvantage of describing state 1 in terms of actions rather than its static state; perhaps you're right that another head of queue example would help.

Response:

The additional example proposed in the response to Intellistor comment 005 is intended to address the confusion. Rather than add more text, a combination of text and graphics should be used to explain the behavior.

#002 (T) Clause 8.6.2, Task Management Example 2 -- HEAD OF QUEUE task

Perhaps if you entered 3 and 7 into the `_bottom_` of the set of enabled tasks, rather than the `_top_`, it would be a little clearer that it simply expedites joining the set of enabled tasks.

Response:

Comment rejected.

Head of Queue tasks should not be shown at some position other than towards the top of the list.

#003 (T) Clause 8.6.2, Task Management Example 2 -- HEAD OF QUEUE task

One other objection I have to the state diagrams is that you show transitions such as "task 2 ended, task 7 created" and "tasks 7, 4 and 6 ended". This implies that a single state change can represent changes to multiple tasks, when in truth this is a `_series_` of state changes.

Response:

Comment accepted.

The descriptions will be modified to portray the state changes as described in the comment.

#003 (T) Clause 8.6.2, Task Management Example 2 -- HEAD OF QUEUE task

...as somebody else pointed out [see Intellistor comment 005], unless you're restricting the "enabled" set to exactly one task, "Head of Queue" is really a misnomer.

Response:

Comment rejected.

The Head of Queue task attribute refers to the task's placement in the queue when the task was created. In that sense, Head of Queue is not a misnomer. Although other terms, such as High Priority, may be better, at this point it is more important to stabilize the terminology.

#004 General

SAM defines a Target Identifier to be 64 bits. However, if we wish to run SCSI over IP (leaving aside for the moment the debate over whether this is a good idea, It `_is_` covered in the GPP draft), 64 bits will be an inadequate address when IPv6 (currently in the design phase) is deployed. IPv6 addresses are 128 bits. I suppose it's too late to change this, but it should be noted.

Authentication is not mentioned anywhere. How are devices to know that they are receiving commands from initiators that they `_should_` listen to? Over SPI, not a problem, since connection via a physical SCSI bus implies ownership of the device, but over a network, it's problematic.

Again in the case of networked SCSI (FC, SSA, GPP, what have you) resource discovery is also a significant problem. Over SPI, traditionally the host attempts to select all devices to establish what devices are present, but over other media this may not be possible.

These last two issues are so big that I don't believe that they haven't been addressed (in fact, I vaguely recall some discussion on the resource discovery problem a long time ago); forgive me if I'm beating a dead horse.

Response:

Defer to SAM-2.

All of the issues you describe must be addressed by the LLP. For example, the target address limitation you mention, could be handled by an LLP function that converts the 64-bit target identifier to an IPv6 address.

The overarching issue here is that there are a lot of implicit LLP functions that are outside the scope of SAM. While a discussion of such issues would serve as a useful guide to implementors and protocol designers, none of these really effect the essential characteristics of the model. For that reason, an in-depth discussion of these matters should be deferred to SAM -2. For now, a few cursory remarks in SAM are in order however.