

T10/07-105 revision 0

Date: March 05, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAM-4: Comments from Adam Richter

1 Overview

The following was received in an E-mail dated 2/28/2007 from Adam Richter. My responses are included as specific changes to SAM-4 for each item (think letter ballot response).

2 Comments

To whom it may concern,

I would like submit some minor language mistakes or opportunities for concision from the SCSI Architectural Model version

4, draft 9.

I am not a SCSI expert or formal participant in any standards body, just a random software engineer who thought it would be educational to try to proofread the latest SAM-4 draft end to end, since your organization has been kind enough to make these drafts freely available.

If you could either forward this note as appropriate, direct me to the correct procedure for submitting these sorts of comments, or otherwise process these comments, I would appreciate it.

I have also appended a note about avoiding the verb "comprise", which appears nine times in SAM-4 draft 9. The advice is applicable to any document written these days that is truly intended to be clear.

To help identify the text that I am referring to, I begin each comment with lines consisting of the relevant section numbers and section names. Suggested new text is in lines that begin with "|".

With the exception of the note about "comprise", all of the other comments are in the order that the excerpts that they refer to appear in SAM-4 draft 9.

2.1 Comment 01

I'll begin with three related excerpts:

4.5.12 Task class

4.5.12.2 CDB attribute

"For a description of the CDB attribute contains a CDB as defined in 5.2 and SPC-3."

4.5.13 Task management function class

4.5.13.2 Task management request attribute

"For a description of the task management request attribute contains a task management request as defined in clause 7."

4.5.13.3 CDB attribute

"For a description of the CDB attribute contains a CDB as defined in 5.2 and SPC-3."

The above three excerpts are each unparseable in the same way. What is the subject of the verb "contains" in each of these sentences? What verb does the preposition "for" modify? Probably these passages are mistranslations from another language.

I suspect 4.5.13.3 should be deleted, since its title and contents are identical those of 4.5.12.2, and "CDB" only appears once in Figure 19 ("SCSI initiator device class diagram"), as an element of "task", not a an element of "task management function."

As for 4.5.12.2 and 4.5.13.2, here are my best guesses for new wording:

| 4.5.12.2 CDB attribute

| "See 5.2 and SPC-3."

| 4.5.13 Task management function class

| 4.5.13.2 Task management request attribute

| "Task management requests are defined in clause 7."

2.2 Response to 01 (editorial)

Change:

"For a description of the CDB attribute contains a CDB as defined in 5.2 and SPC-3."

to:

The CDB attribute contains a CDB (see 5.2 and SPC-3).

Change:

"For a description of the task management request attribute contains a task management request as defined in clause 7."

to:

The task management request attribute contains a task management request (see clause 7).

Delete section 4.5.13.3 as it is a duplicate of 4.5.12.2.

2.3 Comment 02

4.5.15 Level 1 hierarchical logical unit class

In figure 21, page 36 ("Level 1 hierarchical logical unit class"), about two thirds of the way down, in the middle of the diagram is a box with a parenthesized note, which reads:

"{Each instance of a level 1 hierarchical logical unit class, level 2 hierarchical logical unit class, level 3 hierarchical logical unit class, and hierarchical logical unit class 4 hierarchical logical unit class shall contain at least one logical unit or at least one well know logical unit}"

Maybe I misunderstand this note. Why is the parallel construction for unit class 4 in this sentence different from that for 1, 2 and 3?

If there is no reason, perhaps the comment could be shortened to the following:

"Each instance of a hierarchical logical unit class of level 1, 2, 3 or 4 shall contain at least one logical unit or at least one well know logical unit."

2.4 Response to 02 (editorial)

Reject: The wording in the constraint (as it's called in UML) is precise, if wordy. Using the wording you suggest would make it less clear as to what the numbers 2,3, and 4 are. Also, there is no differentiation between the different levels of the hierarchy they all have to have "at least one logical unit or at least one well know logical unit". Also, that wording style is not allowed in ISO standards.

2.5 Comment 03

Also, I am unclear on the use of curly braces in the comment boxes Section 3.6.1 ("Class diagram conventions") has an example of a comment box uses curly braces in Figure 5 on page 17, and then comment box in Figure 6 on page 18 that does not use curly braces. Did I miss an explanation somewhere about what information is conveyed by the presence of absence of curly braces conveys in these comment boxes?

2.6 Response to 03 (editorial)

The curly braces define a constraint on the class or object. If there are no curly braces then the comment is a note. I will add some wording to the model section to explain this more clearly.

Add the following wording onto the conventions section of SAM-4:

A constraint is specified in a class diagram or an object diagram as text encapsulated with a { } notation within a box. See figure xxx for an example of a constraint.

A note is specified in a class diagram or an object diagram as text within a box. See figure xxx for an example of a note.

2.7 Comment 04

4.5.19.2 Logical unit number attribute

The third paragraph is a list, which begins "Logical unit number(s), required as follows:". do not understand the meaning of these words. Perhaps the text could be replaced with the following:

SCSI Devices must define the following logical unit numbers:".

2.8 Response to 04 (editorial)

Change this:

Logical unit number(s), required as follows:

to:

Logical unit number(s), are required as follows:

2.9 Comment 05

4.5.19.3 Logical unit name attribute

"Logical unit name(s), required as follows:". This phrase is unparseable in the same way as the previous one. I think the intended meaning was the following:

SCSI Devices must define the following logical unit names:

2.10 Response to 05 (editorial)

Change this:

Logical unit number(s), required as follows:

to:

Logical unit number(s), are required as follows:

2.11 Comment 06

4.5.23.1 Task class overview

Paragraph 1, sentence 2 has a typographical error: "stated" should be "started." Here is the corrected version:

There shall be one task class for each task that the device server has not started processing.

2.12 Response to 06 (editorial)

Accept as stated.

2.13 Comment 07

5.1. The Execute Command Procedure Call

Output Arguments: Data-In Buffer. last two sentences.

> While some valid data may be present for other values of status, the application client should rely on additional information from the logical unit (.e.g., sense data) to determine the state of the buffer contents. If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider this argument to be unidentified.

I had to do a double take when reading this sentence, because there are a few different possible antecedents to the pronoun "this argument": status, sense data, the service response, and "the buffer contents" (the data-in buffer, that is). I infer that "the data-in buffer" was the intended meaning. In this case, I think that the use of a pronoun has made the sentence a slower read rather than a faster one, thereby defeating the purpose of using a pronoun in the first place. So, I suggest replacing the pronoun "this argument" with "the data-in buffer" for faster reading and, arguably, for clarity, like so:

If the command ends with a service response of SERVICE DELIVERY OR TARGET FAILURE, the application client shall consider the data-in argument to be unidentified.

2.14 Response to 07 (editorial)

Accept as stated.

2.15 Comment 08

5.6.3 - When a SCSI initiator port aborts tasks received on other I_T nexuses

Third paragraph, the parenthesized examples contain an extra comma before the "or", since only two examples were given. That is, "(e.g., COMMANDS CLEARED BY ANOTHER INITIATOR, or COMMANDS CLEARED BY POWER LOSS NOTIFICATION)" should be replaced with the following:

(e.g., COMMANDS CLEARED BY ANOTHER INITIATOR or COMMANDS CLEARED BY POWER LOSS NOTIFICATION)

2.16 Response to 08 (editorial)

Accept as stated.

2.17 Comment 09

8.7 Task priority

The last sentence of the section reads, "However, processing of a collection of tasks with different task priorities should cause the subset of tasks with the higher task priorities to return status sooner in aggregate than the same subset would if the same collection of tasks were submitted under the same conditions but with all task priorities being equal."

"sooner in aggregate" should be "no later in aggregate", since it is possible that exactly the same amount of time would elapse. Also, I wish I could think of a clearer term than "in aggregate" to refer to the completion of the last task and still be reasonably concise, but I haven't been to, so here is suggested text to replace just the word "sooner":

However, processing of a collection of tasks with different task priorities should cause the subset of tasks with the higher task priorities to return status no later in aggregate than the same subset would if the same collection of tasks were submitted under the same conditions but with all task priorities being equal.

2.18 Response to 09 (editorial)

Reject. To me sooner seems clearer. The exact timing is not critical as it is a should.

2.19 Comment 10

8.8 Task state transitions

Below figure 44 ("Task states"), the note for "Translation S2:S4" says, "[...] This is the only state transition that applies to an ACA task". I understand the point that it is making, and it helped my understanding, but this note is technically not quite true. ACA tasks also experience the S0:S2 transition (from "New task accepted" to "Enabled"). So, I guess I would rephrase it as follows:

This and S0:S2 are the only transitions that apply to an ACA task.

2.20 Response to 10 (editorial)

This:

This is the only state transition that applies to an ACA task.

should be:

This is the only state transition out of S2:Enabled that applies to an ACA task

2.21 Comment 11

8.9.3 Ordered tasks

The last sentence of the section has a pretty wordy ending, "Task 5 is required to remain in the dormant task state until task 3 and task 4 tasks complete." I think "task 3 and task 4 tasks" can be shrunk to "tasks 3 and 4", like so:

Task 5 is required to remain in the dormant task state until tasks 3 and 4 complete.

2.22 Response to 11 (editorial)

Reject, that wording style is not allowed in ISO standards.

2.23 Comment 12

I think the verb "comprise" normally should avoided in all avoided in all English language technical specifications, because it has two converse meanings.

The older and more official meaning is the whole "comprises" its parts, but a newer usage, which is more common today, and which I think arose just from people repeatedly making the same mistake, is that the parts "comprise" the whole. All nine instances in SCSI SAM-4 draft 9 appear to be of the second, less officially accepted usage.

Using such a word in technical standards can lead to unnoticed ambiguity or at least delay careful readers.

Addendum - Warning on the use of the word "comprise."

The verb "comprise" appears in the following places in SCSI SAM-4 draft 9:

3.1.22 completed task: "The actual events COMPRISING the TASK COMPLETE response are SCSI transport protocol specific."

3.1.44 function complete: The events COMPRISING this response are SCSI transport protocol specific.

3.5 Numeric conventions

A binary number is represented in this standard by any sequence of digits COMPRISED of only Western-Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b).

[...]

A hexadecimal number is represented in this standard by any sequence of digits COMPRISED OF only the Western-Arabic numerals 0 through 9 and/or the upper-case English letters A through F immediately followed by a lower-case h (e.g., FA23h).

[...]

A decimal number is represented in this standard by any sequence of digits COMPRISED of only the Western-Arabic numerals 0 through 9 [...]

4.4 The SCSI structural model

"The SCSI structural model represents a view of the classes COMPRISING a SCSI I/O system [...]"

4.9 The SCSI model for distributed communications.

"The layers COMPRISING this model and the specifications defining the functionality of each layer [...]"

Interconnect layer: COMPRISED of the services, [...]

7 Task management functions

7.1 Introduction

"Each of SCSI transport protocol shall define events COMPRISING each of these service responses."

How about "corresponding to"?

2.24 Response to 11 (editorial)

Change all nine indicated instances of the term comprising to corresponding to.