

CONGRUENT SOFTWARE, INC.
3998 Whittle Avenue
Oakland, CA 94602
(510) 531-5472
(510) 531-2942 FAX

FROM: Peter Johansson
 TO: NCITS T10
 DATE: March 21, 1998
 RE: SBP-2 Revision 3 Comment Resolution

In the hope of eliminating the tedium of a face-to-face review of editorial changes to SBP-2, I have prepared SBP-2 Revision 3a which includes *almost* all changes referenced below (minor changes agreed in San Diego are not yet incorporated into the draft standard). Responses to technical comments are addressed in a separate memorandum, 98-117.

The prior revision of the document was discussed at the SBP-2 *ad hoc* working group meeting in San Diego on March 17. Some changes were recommended and are incorporated with this revision. With these changes, the working group recommends to the T10 plenary that this document be accepted as resolution of the referenced editorial comments.

Larry Lamers [Adaptec]

Neither of the comments submitted is editorial; they will be addressed in a separate document.

Peter Johansson [Congruent Software]

Technical comments are held for a separate document and all editorial comments save one are accepted.

Reference	Ballot comment	Rejected because
p. 6, 3.1.2.10	Change the definition of "listener" to: A node that receives stream packet(s) for a channel. A listener that receives isochronous data typically observes a single stream packet each isochronous period. There may be zero, one or more listeners for any given channel. The revised definition matches the P1394a changes with respect to asynchronous streams and "loose" isochronous.	While this is strictly correct, SBP-2 makes no use of asynchronous streams; it is likely more informative to the reader to leave the definitions of "listener" and "talker" as they are.

Gene Milligan [Seagate Technology]

Comments 1 – 10, 13, 15, 19, 22 – 24, 27, 28, 31 – 35, 37 – 42, 45, 47, 48, 51 – 56, 58, 59, 65, 66, 68, 72 – 81, 83 and 84 are considered editorial and accepted. Other comments, although editorial in nature, are rejected for the reasons given below. Comments 25, 26, 43, 61, 63, 64 and 69 – 71 are either technical or controversial and will be addressed in a separate document.

Reference	Ballot comment	Rejected because
11	The membership list for NCITS looks like a remarkably old X3 list.	A correct observation, but I reckon it'll give the ANSI editor(s) something to do when they get the document.
12	Delete the revision history but it would be nice to keep it in a separate document.	Pete McLean made the suggestion that the revision history be moved to be part of the "front material" that is ultimately deleted by the ANSI editor before publication. This change will be made in a future draft.
14	I suggest changing the purpose towards the 1394 not just a serial Interface which was previously accomplished and moving all the desired information to the foreword.	I think the design goals of SBP-2 merit placement here instead of a foreword.
16	Is there a reason not to list the published SAM and other command sets?	I believe that SBP-2 implementations cannot comply with SAM or SPC, hence they are not referenced.
17	It would be better to include the key words listed under conformance in a Key Word clause as is done in most of the other SCSI standards.	SBP-2 is not a SCSI standard. Because most of its expected readers are assumed to be familiar with IEEE Std 1394-1995, it adopts many of its editorial cues from styles used in IEEE standards. The use of a conformance glossary is consistent with ANSI style guides.
18	Expected should be compared to the other SCSI standards to determine why the definition should be different.	The definition of "expected" conforms to usage in IEEE Std 1394-1995 and in any case does not differ substantively from the definition used in other NCITS standards.
20	Regarding 3.1.2.1 are bytes only used for data?	Yes—but it's hard to imagine anything precluded by the generic "data."
21	Regarding 3.1.2.2 there seem to be extraneous 's as there are in 3.1.2.34.	If the phrase appeared by itself it would be "command set-dependent", hence the presence of both hyphens.
29	What does 'The meaning of the field shall be defined by the organization responsible for the unit architecture' mean? How is it different than vendor dependent (which in other SCSI standards is vendor specific)?	In the CSR architecture; the specifier of the unit architecture and the vendor are not necessarily the same. For instance, SBP-2 is a unit architecture specified by NCITS but within that unit architecture room has been left for different specifiers of command sets. The use of "-dependent" in place of "-specific" is conventional in the IEEE standards upon which SBP-2 is based.

Reference	Ballot comment	Rejected because
30	I assume the use of terms differing from other members of the SCSI family of standards is not NIH but is congruence with 1394 and/or 13213. If the assumption is correct, it would be helpful to add an informative annex providing a correlation table of 1394/SCSI terms (e.g. lsb/LSB; vendor dependent/vendor specific). If I am wrong and it is NIH then change the terms to the SCSI family terms.	It is not a case of “not invented here” but a matter of editorial conventions familiar to the expected audience. SBP-2 is not a SCSI standard, it is a standard produced by T10, Technical Committee on Low-level Interfaces. I believe the correlation between vendor-dependent and vendor-specific is self-evident.
36	With regard to figure 6, is the ground symbol at the end of a linked list conforming to some standard for illustrating linked lists? To me it seems odd.	I don't recall when I started employing this convention or whether I adopted it or made it up out of whole cloth. At any rate, I think its usage is clear even if quirky.
44	The break lines (I presume) in Figure 14 are peculiar.	They look OK to me.
46	Regarding 'A target's command set and device type determine the length of the ORB's which shall be fixed for a particular command set and device type.' In 5.1.2, some command sets have different length command blocks. Should this be 'A target's command set maximum length and device type determine the length of the ORB's which shall be fixed for a particular command set and device type.'? If so is padding defined somewhere?	Padding is outside the scope of SBP-2 and is left to command set documents that define how SBP-2 is to be used. For example, Annex B shows how a SCSI CDB is to be contained within an SBP-2 ORB. I think the other language, about how the target's command set and device type determine the ORB size, is clear.
49	Regarding figure 16, how do you determine the length of the command block since the device seems to be able to choose a fixed length?	This is beyond the scope of SBP-2 and is the province of the relevant command set standard.
50	The unlabeled table does not seem to match the numbering convention. It seems odd the 0-8 have no subscripts but F does.	Zero through 8 require no subscript to be disambiguated and hence have none. Although I suppose one could argue that the single digit F needs no subscript either... In any case, I've added the following to the clause on numeric conventions: “When the subscript is unnecessary to disambiguate the base of the number it may be omitted.”
57	5.3 is redundant to the requirements of 5.1. It would be better if 5.1 deleted the requirements in favor of 5.3.	There is some redundancy here but I don't think it harmful. I think the description of the <i>notify</i> bit belongs in 5.1. If I had to eliminate anything I would modify 5.3 but prefer to leave both sections as they are.
60	I think Std may be a standard abbreviation for standard but I do not think it is for ANSI. I suggest that a 'Std' be globally replaced with 'standard'.	In the context used it is part of the formal designation of the standard, IEEE Std 1394-1995.

Reference	Ballot comment	Rejected because
62	In 7 shouldn't 'all other directories and leaves' be 'all other directories and leafs'?	No, my dictionary lists "leaves" as the plural of "leaf."
67	Larger and/or darker font in figure 58 would be helpful.	The over-arching concern was to fit the state diagram on one page while still maintaining legibility. The font sizes and weights are identical to those used in IEEE Std 1394-1995.
82	Regarding 'F.8 Hard reset A Serial Bus reset causes the target to execute a hard reset, as defined by SAM-2.' Does Annex A 'A Serial Bus reset, by itself, shall not alter a target's responsiveness to request subactions.' match the F.8 claim?	I believe it does, "Request subactions" is a technical term within IEEE Std 1394-1995 that refers to the Serial Bus primitives read, write, lock, etc. I don't think SAM hard reset behavior affects how these are handled.

Steve Finch [Silicon Systems]

Comments 3, 10 – 14, 16, 17, 19, 20 – 23, 25, 26, 28, 29, 31 – 33, 35 – 38, 40 – 41, 45, 46, 48, 49, 51, 55, 61, 62, 64, 66, 70, 74, 75, 80, 81, 83, 84, 87 – 91, 94, 95, 96, 104, 107, 109, 112, 116 – 118, 130, 137, 139, 140, 141, 143, 144, 150, 152, 154, 155, 157, 163, 165 – 169, 174, 177 – 179, 182, 185 and 187 are accepted as editorial. Other comments, although editorial in nature, are rejected for the reasons given below. Comments 1, 4, 5, 6, 7, 8, 18, 34, 43, 44, 47, 50, 53, 56 – 60, 63, 65, 67 – 69, 71, 72, 73, 76 – 78, 85, 86, 92, 98 – 102, 105, 106, 108, 110, 111, 113, 114, 119, 121 – 129, 131 – 136, 142, 145 – 149, 153, 158, 159, 161, 162, 170, 171, 175, 176, 180, 181 and 186 are either technical or controversial and will be addressed in a separate document.

Reference	Ballot comment	Rejected because
2	Remove the text title "Revision history" on pages ix through xviii. This information is not appropriate for inclusion in a standard.	See response to Gene Milligan's comment 12.
9	Section 3.1.2 There are a lot of related and/or similarly named terms such as command block, normal command block, stream command block, command descriptor block, stream command ORB, stream command block ORB, and other combinations of these terms with other terms such as agent. I don't believe these are consistently used and nor adequately defined. I propose that we limit the number of such terms to the following list, define these terms in section 3.1.2, and review the document to make sure the usage of these terms is consistent.	The draft has been reviewed on more than one occasion for consistency; corrections of particular errors are welcome. I do not believe that the proposed definitions add value.
15	Section 3.1.2.25, change "stream ID" to "stream identifier (stream ID)".	Accepted, but in an inverse fashion: all the instances of "stream identifier" have been changed to "stream ID."

Reference	Ballot comment	Rejected because
24	Section 4.4, last paragraph, delete "(or a scatter/gather list)". This terminology is used only twice, each time as parenthetical adjunct to insert a second name for an unrestricted page table. This second name is never used elsewhere, so there is no reason to include it at all. The other location is in section 5.1.2.1.	The phrase "scatter/gather list" is used more commonly to describe what SBP-2 formally names an "unrestricted page table". I believe it adds value.
27	Section 4.5, last paragraph, last sentence, delete "The time critical nature of isochronous operations requires that". This is not necessarily true, depending upon the isochronous control ORB's that are issues and the functionality supported by a target device.	This statement is true of some isochronous control functions. If linked lists were not supported there would be inadequate time for the initiator to receive completion notification for one request and then signal the next.
30	Figure 9 has gray scale material that leads to understanding of the figure. When I printed out the .pdf file, the differences in the grayscales is marginal at best. I suggest that the gray scale be changed to some type of fill pattern to aid in the reliable production of the standard.	I will discuss this with the ANSI editor as the draft moves to its final stages; I experimented with fill patterns and found the figure less digestible.
39	Section 5.1.2, delete the NOTE. It adds nothing to the standard and makes a statement which may or may not be true.	There was considerable discussion over what to keep in the ORB header (the first 20 bytes) in order to achieve a workable 32-byte ORB. I think this note captures a smidgen of that history and should stay.
42	Section 5.1.2.1, fourth paragraph after Table 1, delete "(also known as a scatter/gather list)". This terminology is used only twice, each time as parenthetical adjunct to insert a second name for an unrestricted page table. This second name is never used elsewhere, so there is no reason to include it at all. The other location is in section 4.4.	See response to comment 24.
52	Section 5.1.3, delete the NOTE after Figure 19 or clarify. If you want to include it, I can come up with a thousand more "usage" notes. Why add this one only?	The point of the note is not the example given (one could find many other examples) but to provide at least one example as to why a mapping of two or more source channels into a single destination channel can be meaningful.

Reference	Ballot comment	Rejected because
54	Section 5.1.3: Add a table that shows which fields are valid for which command and, in appropriate cases, what values within a field are valid on a per command basis. For example, the stream_event field is only valid for START, STOP, PAUSE or UPDATE CHANNEL MASK commands, and the values that are valid for these commands depend on whether the stream is a listener or a talker.	I think this information is already accurately and clearly presented.
79	Section 5.2, second paragraph, why is via in italics?	Common editorial practice.
82	Section 5.2.2, add the following to the last sentence of the NOTE: "i.e., page_size is equal to 4."	I think the caption on Figure 31 says it all.
93	Section 7.6 shows the first quadlet of the leaf as having a size and CRC. This quadlet exists for the unit directory and the logical unit director but sections 7.4 (unit directory) and 7.5 (logical unit directory) do not mention these fields.	Neither the unit directory nor the logical unit directory has an illustration; prior familiarity with the CSR architecture is a prerequisite for reading this section.
97	Section 8.2.1, change the Note to read: "NOTE - The speed at which the block write request to the MANAGEMENT_AGENT register is received shall determine the speed used by the target for: - all subsequent requests to read the EUI-64 from the initiator's configuration ROM - all subsequent requests to read ORB's from initiator memory, or store status at the initiator's status_FIFO associated with the logical unit once access to the logical unit has been granted to the initiator. Command block ORB's separately specify the speed for requests addressed to the data buffer or page table."	The existing note is accurate and not improved by the proposed text.
103	Section 8.2.1, an observation: This section (and similarly section 8.2.2 in regards to create streams) covers some of the actions of both the initiator and the target. For example, the text says that the initiator signals the login ORB to the target, and we state that the target reads the EUI-64 and returns a login response, but we never say the target reads the ORB nor that the target returns a status. I feel that we should make this be a more definitive list of the actions of both side.	The actions of reading an ORB and returning status are hard to avoid for any ORB; they are not particular to these requests, hence are not described here. Similarly, an initiator must signal an ORB to a target else the target remains unaware of it; this is also generic and not recapitulated here.

Reference	Ballot comment	Rejected because
115	Section 9.1, second paragraph, first sentence, change "The other target agents, command block and stream control," to "The other target agents, normal command block, stream command block and stream control".	The proposed text seems verbose to me and doesn't add information not already present.
120	Section 9.1.3, change first paragraph to read: "Any initiator application that executes in a single-threaded environment, such as BIOS, has little need of the target fetch agent's capabilities to manage multiple outstanding requests. Such applications may use a simpler procedure than that described in 9.1.2 to signal requests to the target. Subsequent to initialization of the target fetch agent by means of a write to the AGENT_RESET register, the BIOS may signal one request at a time to the target as follows:"	This whole (informative) clause is about the BIOS; the reader is left to discover that it is equally applicable to "any similar application." The attempt to make it seem more generic breaks down in the last sentence of the paragraph at the left, where an abrupt and unexplained transition is made from "any initiator application" to "the BIOS."
138	Section 10.4.1, first paragraph, second sentence, change "targets may also recognize task management ORB's to abort tasks." to "targets may also recognize ABORT TASK ORB's."	This is not consistent with the editorial style I have been attempting.
151	Section 11 should be reorganized as follows: 11. Intro to recorded information and relationship of Data Format Records and cycle mark index entries. Explanation that, in the case of CIP format, the internal format of isochronous data packets must be understood in order to process the stream.	This is essentially the organization of section 11 today, without distinguishing the introductory material and table with its own clause number.
156	Section 11.2, in the first paragraph below the table, change "may have been transformed" to "shall be transformed". It is transformed even if the transformation causes no change in values.	The statement is correct as is. A target does not necessarily know whether or not channel transformations were in effect at the time of recording. The identity mapping is not considered a mapping in common-sense usage.
160	Section 11.5, figure 64: change the field "fmt-dependent" and any textual references to "fmt_dependent".	The description "fmt-dependent" is not a name of a field (as can be seen by its use in two figures).
164	Section 11.5, figure 66, names a field the same name as used in figure 64. This is confusing and could easily be changed to another name.	See response to comment 160.
183	Section A.2, in the paragraph starting "Once a target", change "target" to "node". A target is equivalent to a unit, not a node.	Strictly speaking, this is correct as written—a target is responsible for its own address ranges within configuration ROM. But this text applies equally to initiators and has been moved to the start of the annex (with the suggested wording change).

Reference Ballot comment

184 Section B.2, in the first paragraph after the second table, change "shall specify the content" to "shall specify the validity of the content".

Rejected because

This matches the wording used in SPC to describe the analogous field in SCSI sense data.

Jeff Williams [Western Digital]

Comments 1 – 9, 11 – 13 and 15 – 17 are accepted as editorial. Comments 10 and 14 are either technical or controversial and will be addressed in a separate document.