

DocNum: T10/01-265r0  
Author: Eric Anderson  
Title: Minutes of SBP-3 Working Group  
ftp://ftp.t10.org/t10/document.01/01-265r0.pdf

**Minutes of the SBP-3 Working Group meeting, August 22-23, 2001**  
Apple Computer, Cupertino

Attendees:

Eric Anderson	Apple	ewa@apple.com
Firooz Farhoomand	Panasonic	firoozf@ix.netcom.com
Lee Farrell	Canon	lfarrell@cissc.canon.com
John Fuller	Sony	jfuller@computer.org
Andy Green	Oxford Semiconductor	andy.green@oxsemi.com
Peter Johansson	Congruent Software	Pjohansson@ACM.org
Steve Powers	Microsoft	spowers@microsoft.com
Mike Teener	Zayante	mike@zayante.com

The following agenda was presented by Johansson. In the minutes that follow, the start of discussion of items listed below is denoted by the index number listed within square brackets, such as [4.1]. Note that these references do not always appear in order, and may not signify the conclusion of discussion of a previous agenda item.

1. Introductions and procedures
  - 1.1 T10 Membership and voting
  - 1.2 Document naming conventions
  - 1.3 Two-week rule
  - 1.4 Meeting fees
  - 1.5 Approval of prior minutes
2. Call for patents
3. Informal liaison
  - 3.1 IEEE P1394.1 [Johansson]
  - 3.2 IEEE P1394.3 [Johansson]
4. Prior action items
  - 4.1 P1212 Revision entry in unit directory [Johansson]
  - 4.2 Request AV/C expert to review Annex H [Fuller]
  - 4.3 Request AV/C expert to define track metadata [Fuller]
  - 4.4 Operational description of login (bus reset) [Johansson]
  - 4.5 SBP-3 Liaison report to 1394 TA AVWG [Anderson]

5. Review of changes in working draft
  6. Old business
    - 6.1 New isochronous model
  7. New business
    - 7.1 FAST\_START enhancements [Green / Johansson]
    - 7.2 Efficient use of the page table [Rice]
    - 7.3 Stream command block ORB [Johansson]
    - 7.4 Annex H
    - 7.5 Cycle time transforms and tolerance [Green]
    - 7.6 Bridge scenarios
  8. Meeting schedule
  9. Review of action items
- Adjournment

[1] Johansson called the meeting to order and updated the agenda, as reflected above.

[1.3] Johansson briefly reviewed the two-week rule, explaining that it did not prevent the discussion of documents posted less than two weeks before a meeting.

[2] Johansson reviewed general T10 policies and procedures. In general, attendance and participation at T10 ad hoc meetings (such as this one) is open to both visitors and T10 members. When formal votes are taken, either in an ad hoc meeting or in the T10 plenary, one vote is permitted each organization, to be cast by its principal representative or designated alternative. A two-week rule is in effect: No matter may be voted on unless notice was given at least two weeks prior. Documents to be voted on must have been posted two weeks prior to the vote. The two-week rule can be waived if nobody objects. Announcements of new documents and meetings must be posted to the T10 email reflector; all other business can be conducted on the working group reflector.

The following paragraph about ANSI/T10 patent policy is copied from past T10 Plenary minutes:

A document is available from ANSI, "Procedures for the Development and Coordination of American National Standards", at no charge. This document is also on the web at [http://www.ncits.org/help/ansi\\_sdo.html](http://www.ncits.org/help/ansi_sdo.html). Section 1.2.11 contains the ANSI patent policy. Amy Marasco manages patent issues for ANSI and can be contacted at [amarasco@ansi.org](mailto:amarasco@ansi.org) or 212-642-4954. Gene Milligan prepared a useful "Handy dandy Technical Committee's Patents Guide", which is available at <ftp://ftp.t10.org/t10/document.99/99-291r0.pdf>.

[1.5] The minutes from July 17 (Colorado Springs) were approved without objection:

<ftp://ftp.t10.org/t10/document.01/01-233r0.pdf>

[3.1] Johansson reported that the 1394.1 Comment Registry preparation is nearly complete, adding that comments had been recirculated to balloters for confirmation. Johansson said the registry will be published in a few weeks, then the BRC will attempt to resolve all 508 comments. The BRC meeting(s) are not yet scheduled, and it is likely that easy comments may be resolved by email in advance.

[3.2] Johansson noted that the Alan Berkema (the vice chair) is coordinating the 1394.3 BRC, for which the comment registry has been prepared but not yet posted. Johansson added that the situation was similar to that of 1394.1.

[4.1] Johansson reminded everyone that SBP-3 had requested a Revision Entry key in IEEE 1212r, adding that he doesn't know if there is any controversy regarding the request yet.

[4.2] Fuller stated that he has made some progress in mastering AV/C, but more work remains.

[4.3] Fuller reported that this work is not yet done, because he is still working on item 4.2.

[4.4] Johansson commented that he has not yet completed this item.

[4.5] Anderson said that he had given this liaison report to the 1394 Trade Association Audio Video Working Group and had subsequently posted a URL for the recent minutes of SBP-3 on the AVWG reflector. Johansson asked what could be done to get the AVWG more involved. Fuller noted that AVWG often works by designing a product and then publishing a standard to describe its function, so our approach of writing a standard first may be unfamiliar to some of that group's members.

The transport of FCP (Function Control Protocol, from ISO/IEC 61883 part 1) over SBP was discussed [relates to 4.2]. Johansson pointed out that an AVC command execution engine could process the same commands delivered by SBP without needing significantly more resources. Anderson noted that if a target fetches FCP frames through ORBs, the frames will arrive as response packets, and the target would have greater control over the arrival scheduling than with conventional FCP, where block writes containing requests can be received at any time.

[5] Johansson led a review of changes in the latest draft of SBP-3:

<ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01e.pdf>

Everyone agreed with the new text added to section 5.2 (page tables).

Johansson reviewed section 7.6 (config ROM directory entries) and its subsections, containing new edits from the previous meeting, and noted that he had cleaned up imprecise uses of "logical unit" and "target" (such as in 7.6.11). Johansson also called attention to a new usage note in section 7.6.7 and the new "d" bit in section 7.6.9.

Johansson concluded the draft review by noting that section 11.4 (cycle mark index) had been removed.

[7.1] Johansson led a discussion of updated Fast Start documentation:

<ftp://ftp.t10.org/t10/document.01/01-248r1.pdf>

Anderson suggested changing "(available to the target as a result of login)" to "(available to the target as a result of login or reconnect)" after Figure 46, and in a later paragraph.

The final paragraph of section 6.4.7 was modified to only accept fast start packets when in the RESET state if prev\_ORB is null; if prev\_ORB is not null the packet shall be ignored (and not compared with ORB\_Pointer). Various other touchups made in this section.

In figure 68 and the corresponding text, transition F0:F3 was updated to happen only if prev\_ORB is null; not if prev\_ORB is equal to ORB\_POINTER.

Anderson moved to roll 01-248r1 into the draft as modified.

Fuller seconded.

Motion passed with none opposed.

[7.2] Anderson collected possible topics for a proposed annex on performance and implementation tips, as follows:

1. Use 2 kilobyte payload packets, or whatever the maximum is based on your PHY speed.
2. Use the fastest PHY available.

3. Read page table entries only once, in advance, preferably as a block.
4. Support Fast Start and use it at every opportunity.
5. Use `ack_complete` wherever allowed, instead of `ack_pending`.
6. Use physical memory on the initiator, where supported - for ORBs, page tables, etc..
7. Use ORB chaining if sufficient work exists, rather than letting the target go idle after each I/O.
8. Prefetch an ORB and at least part of its page tables during the previous ORBs execution, if a valid next ORB link is available.
9. 1394.1 bridge-aware targets should pipeline payload transfers: after seeing an `ack_pending` on a block read or write request, issue the next block read or write immediately (if any), because just a single interposed store-and-forward bridge will otherwise cut throughput roughly in half.

[6.1] Johansson presented an updated simplified isochronous model proposal:

<ftp://ftp.t10.org/t10/document.01/01-222r0.pdf>

Johansson explained figure 18 and its many fields, noting that many combinations of fields were meaningless, and that profiles showing typical usage should be added to the document. Johansson commented that if some of the information was also available from plugs, those parameters would not be repeated in the ORBs in such a case.

Anderson agreed with Johansson's suggestion that a byte-granular command set for disk read and write would be most suitable, rather than trying to wrap byte granularity around an existing block-oriented disk access command set such as RBC.

Johansson noted that for some kinds of devices it might not be necessary to use the create stream ORB, as the login might return the required stream, if no other kind of agent exists. This would be specified by the command set used by such a device.

Fuller asked why the 6th quadlet of figure 18 was not command set dependent. Johansson explained that a command set could be ignorant of these 1394 details, as long as the ORB provides a place to specify them, which simplifies

command sets. Johansson also noted that in some implementations, not all of the information in that quadlet would be needed.

Fuller, Johansson, and Teener discussed the architecture at length, including how bandwidth allocation would be performed, and the treatment of redundant fields when used with certain protocols or data formats. Johansson argued that it was best for the ORB to provide space for each parameter even if they would sometimes be redundant, because nothing in 1394 required protocols or command sets to specify these 1394-specific values (channel number, speed, etc.) within their own structures.

[7.4] The transport of FCP within SBP was revisited. Fuller preferred the use of a bidirectional ORB to point to a request and a response (possibly the same buffer), rather than placing FCP frames directly in the ORB. Fuller felt that allowing small FCP frames to go "in line" in the ORB would add complexity, and was not worth the potential efficiency savings. Though the packet count compared to plain FCP would be double or more, this service is not intended for high-performance exchange (ordinary SBP would be used), so the inefficiency would be insignificant. Johansson noted that an FCP frame can be up to 512 bytes, which if wrapped in an ORB would then be too large to transport at s100, so Fuller's proposal provided more flexibility. Fuller also noted that with variable-size commands, a target would have no way to know what size ORB to fetch (with inline commands) and would have to either read the ORB piecemeal, or fetch the maximum possible ORB, which could be very inefficient.

[7.3] Johansson led a discussion of the usage of the various fields in the stream command ORB. Johansson and Fuller discussed how 1394.1 bridges would be involved, with Fuller advocating that the use of Plug Registers to describe isochronous connections would solve the bridging problems. Fuller noted though that Plug registers as presently defined have only two bits available for speed encoding, which does not cover speeds added in IEEE P1394b. It was suggested that SBP-3 could define additional speed bits for plugs. Johansson drew a usage example for isochronous access for a DVD player, highlighting which fields would be used and not used. Johansson noted that the DVD drive could use existing MMC commands plus a 1394 (61883-1) Plug, so that specifying isochronous parameters in the stream command ORBs would be unnecessary. Anderson posited that AVWG didn't appear to be moving quickly to speeds above s100, and AVWG was unlikely to drive the definition of higher-speed plug bits, so the SBP-3 or 1394.1 group could do so with little risk of contention. Johansson noted that with reliance on plugs, SBP-3 could simply specify an "isochronous" bit in the existing ORB structure, and leave all attributes of isochronous transport up to the command set - to resolve with plugs, commands, or whatever as needed.

Johansson summarized what had been proposed: Simply add an isochronous bit to the ORB, keep the "create stream" ORB, add high speed plugs to 1394.1, get rid of the stream command ORB, and rely on plugs or other command-set specific services to establish 1394 isochronous parameters. Anderson noted the need to be able to release a stream without performing a logout.

[7.4] Regarding FCP's ability to store Interim status; Anderson noted that the initiator could use nonphysical memory so that no race condition exists (though this may require more resources). Anderson said that using unsolicited status to store interim status was undesirable because it would bottleneck all transactions on a single resource, much like native FCP does. Anderson preferred Fuller's suggestion to double the size of the response buffer, and use one part for interim status (if any) and the other for final status - assuming there is only one interim status report per ORB.

Anderson suggested allowing SBP ORB status to have a new "interim" flag indicating that the ORB is still in progress, but has generated some kind of status, the meaning of which would be command-set dependent. Johansson said that status with a src value of 3 could convey this, adding that command sets could use either the status block or a data buffer or both.

Johansson suggested that if two responses are targeted at a buffer (one being interim), then the main response should be specified to be at offset zero, and the interim response at offset 512 or some known offset, to be defined in annex H. Anderson agreed that the main response should be at offset zero. Johansson suggested that the interim response offset be specified in the CDB portion of the ORB which is presently unused.

Johansson said he would write a draft showing how isochronous services would be used if we adopted the new "isochronous bit in normal ORBs" plan.

[7.5] The group discussed Neil Buxton's email question about cycle clock accuracy, and Scott Smyers' response. All agreed that Neil's question was not specific to SBP-3, and pertained to any record and playback of 61883-type data. Green gave an explanation of how a media stream transported over 1394 can obey an external (non-1394) clock that is of higher precision than 1394's 100ppm clock, such as NTSC colorburst (30 ppm). To do so, the source generates timestamps based on the external clock, and the sink recreates the external clock. Green explained that an SBP-3 disk drive (or computer) that tried to record and later play back such data would need to participate in the external clock, because using the 1394 clock alone would not be sufficient to generate a properly synchronized stream. Anderson summarized by saying that in any application transporting data through 1394 using a higher-precision non-1394 clock, it would generally be impossible to regenerate such a stream using only

the 1394 clock as a source, because the required accuracy would be lacking. The group concluded that Neil had raised valid issues, though they were not unique to SBP-3, as they would (for example) apply to a conventional computer recording and replaying an isochronous stream. Johansson noted that 1394 did not require the use of higher-precision clocks, so an SBP-3 device could record and play back streams using just the 1394 clock if the stream was clocked with a precision equal to or lower than 1394's 100 ppm tolerance.

[7.6] Johansson led a continuation of past discussions of how to make SBP-3 work with 1394.1 bridges. One topic of interest was how to deal with Net Generation updates. Johansson suggested that since these events are rare, it seemed appropriate to treat them like an SBP-2 bus reset: abort the task set, and require a reconnect. The target and initiator may see the Net Generation update in either order, but packets can't pass a Net Generation update wavefront, so the initiator can start a reconnect immediately and have no danger of getting out of sync with the target (e.g., performing a reconnect before the Net Generation update). The reconnect time could be the same as in SBP-2, with an assumption that a bridge-aware initiator can take the (known or predicted) net diameter and population into account when selecting the reconnect period.

Fuller noted that if an initiator requested a login of a target on the same bus, using a global node ID instead of a local address, things would be very inefficient because a bridge would echo every packet, so it is important for initiator implementations to be smart about this. If 1394.1 restricts the "echo" service, using a Global Node ID on the local bus might not even work. Johansson noted that learning a Global Node ID for a local node is difficult, because DEP (Device Enumeration Protocol) returns local node IDs when possible. Johansson concluded that it might be best to not even mention this possible problem.

Fuller noted that SBP-3 had no way for a target to learn that a remote initiator had vanished, and thereby release its resources for use by other initiators. Anderson pointed out that legacy initiators won't be able to request resolution of this kind of problem, so perhaps bridge-aware targets should revalidate their connections when a legacy login or query logins arrives if there was no recent activity. Johansson noted that login and query logins are so rare that SBP-3 could just always require validation, except perhaps in cases where the login can succeed because one unused login is available. Fuller noted that a remote timeout may be needed to revalidate existing logins, and explained that failing to validate a login would not clear the login, it would instead start the reconnect timer.

Anderson noted that if SBP-3 denies a login to a legacy initiator, that initiator might not ask again. So even though the target has to start reconnect processing when it suspects a login has gone stale, it would be best to try to



complete the reconnect processing before answering the new legacy login request. Anderson noted that legacy software likely honors the Management ORB timeout in the Config ROM (adding that Mac OS does this), and so SBP-3 devices can publish a large Management ORB timeout in order to "buy time" to clear out dead logins and increase the odds of granting a legacy login on the first try.

Someone mentioned that 1394.1 needs to address the use of the busy retry register (offset x210); noting that it resets due to a bus reset, and also noting that an off-bus target's busy retry controls the use of ack\_busy between that target and the closest bridge. Software wanting to set a target's busy\_timeout might want to set the bridge portal's busy\_timeout instead (or might not).

[8] The upcoming meeting schedule is:

October 3 - 4 (Portsmouth, NH)  
November 6 - 7 (Monterey, CA)  
January 21 - 22 (Hawaii)

The December meeting was canceled by consensus.

Adjourned.

-----

General information and document index

The SBP-3 email reflector SBP3@isg.apple.com can be accessed as follows:

Subscribing:

email requests@isg.apple.com w/subject "subscribe sbp3"

Help?:

email requests@isg.apple.com w/subject "help"

An automated system had been created for the allocation of T10 document numbers, and the subsequent submission of documents for posting:

<http://www.t10.org/members/ad.htm>

The following documents have been posted pertaining to SBP-3:

00-328 Eric Anderson  
Fast Start proposal (PowerPoint slides)

<ftp://ftp.t10.org/t10/document.00/00-328r0.pdf>

- 00-371 Peter Johansson  
Minutes of SBP-3 Study Group September 19, 2000  
<ftp://ftp.t10.org/t10/document.00/00-371r0.pdf>
  
- 00-388 Peter Johansson  
SBP-3 Project Proposal  
<ftp://ftp.t10.org/t10/document.00/00-388r0.pdf>
  
- 01-057 Eric Anderson  
Fast Start Proposal  
<ftp://ftp.t10.org/t10/document.01/01-057r0.pdf>
  
- 01-060 Eric Anderson  
Minutes of SBP-3 Working Group January 24-25, 2001  
<ftp://ftp.t10.org/t10/document.01/01-060r0.pdf>
  
- 01-067 Lance Flake  
RBC Access For AV/C Data Interchange  
<ftp://ftp.t10.org/t10/document.01/01-067r0.pdf>  
<ftp://ftp.t10.org/t10/document.01/01-067r1.pdf>
  
- 01-070 Peter Johansson  
Bridge-aware targets and node handles  
<ftp://ftp.t10.org/t10/document.01/01-070r0.pdf>
  
- 01-101 Eric Anderson  
Minutes of SBP-3 Working Group March 6-7, 2001  
<ftp://ftp.t10.org/t10/document.01/01-101r0.pdf>
  
- 01-102 Scott Smyers  
Proposal for modifications to SBP3 and RBC  
<ftp://ftp.t10.org/t10/document.01/01-102r0.pdf>
  
- 01-103 Firooz Farhoomand  
Using SBP-3 for DVD playback  
<ftp://ftp.t10.org/t10/document.01/01-103r0.pdf>
  
- 01-137 Peter Johansson  
Stream command block ORB  
<ftp://ftp.t10.org/t10/document.01/01-137r0.pdf>

- 01-138 Peter Johansson  
Bi-directional ORBs (PowerPoint slides)  
<ftp://ftp.t10.org/t10/document.01/01-138r0.pdf>
- 01-139 Eric Anderson  
Minutes of SBP-3 Working Group April 26-27, 2001  
<ftp://ftp.t10.org/t10/document.01/01-139r0.pdf>
- 01-179 Andy Green  
Proposal to modify isochronous recording format  
<ftp://ftp.t10.org/t10/document.01/01-179r0.pdf>
- 01-180 Peter Johansson  
RBC-2 commands for extent management  
<ftp://ftp.t10.org/t10/document.01/01-180r1.pdf>
- 01-187 Eric Anderson  
Minutes of SBP-3 Working Group June 5-6, 2001  
<ftp://ftp.t10.org/t10/document.01/01-187r0.pdf>
- 01-200 Peter Johansson  
Distributed Buffers  
<ftp://ftp.t10.org/t10/document.01/01-200r0.pdf>
- 01-223 Eric Anderson  
Minutes of SBP-3 Working Group July 17-18, 2001  
<ftp://ftp.t10.org/t10/document.01/01-223r0.pdf>
- 01-248 Peter Johansson  
MP-friendly Fast-Start  
<ftp://ftp.t10.org/t10/document.01/01-248r1.pdf>
- 01-265 Eric Anderson  
Minutes of SBP-3 Working Group August 22-23, 2001  
<ftp://ftp.t10.org/t10/document.01/01-265r0.pdf>
- 01-287 Peter Johansson  
Bare-bones Isochronous  
<ftp://ftp.t10.org/t10/document.01/01-287r0.pdf>

Latest draft SBP-3 document:

<ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01e.pdf>

[end]