

DocNum: T10/02-208r0
Author: Eric Anderson
Title: Minutes of SBP-3 Working Group
ftp://ftp.t10.org/t10/document.02/02-208r0.pdf

Minutes of the SBP-3 Working Group meeting, May 29-30, 2002
Timberline Lodge, Oregon

Attendees:

Eric Anderson	Apple	ewa@apple.com
Robert Botchek	Granite Digital	rbotchek@granitedigital.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

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 Review P1394.1 BRC plans with respect to BUSY_TIMEOUT register [Johansson]
 - 4.2 Update Annex E to include Revision entry and correct value for Version entry [Johansson]
 - 4.3 Publish revised 02-069r1 [Johansson]
5. Review of changes in working draft

6. Review reflector traffic
7. Old business
 - 7.1 AVD Commands
 - 7.2 Bridge-aware SBP-3 target operations
 - 7.3 Update status_FIFO address based on RECONNECT request
 - 7.4 Processor model and messages for extent manager
 - 7.5 Dynamic LUNs
8. New business
9. Meeting schedule
10. Review of action items
11. 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.

[1.5] Anderson noted that he had distributed draft minutes from January 21 (Waikoloa) and March 12 (Dallas) the previous night. After review, the group approved the minutes from both meetings.

<ftp://ftp.t10.org/t10/document.02/02-206r0.pdf>

<ftp://ftp.t10.org/t10/document.02/02-207r0.pdf>

The minutes from November 6 (Monterey) were not yet available for approval.

[1.1] [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. Section 1.2.11 contains the ANSI patent policy. Amy Marasco manages patent issues for ANSI and can be contacted at 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>.

[3.1] Johansson reported that the IEEE 1394.1 BRC had met earlier in the month in Zurich, at which time all but 134 of 517 ballot comments were resolved. Johansson said the BRC plans to conduct a recirculation ballot in 2002. Johansson added that isochronous setup/teardown and a new loop healing scheme were major areas to still be addressed.

[3.2] Johansson noted that activity in IEEE 1394.3 had picked up, and he predicted a recirculation ballot within a month or so.

[4.1] [4.2] [4.3] Johansson stated that his three action items from the previous meeting were complete.

[5] Johansson led a review of changes in SBP-3 draft "g":

<ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01g.pdf>

Johansson remarked that section 4.7 had been updated to remove support for hints, as per the previous meeting.

Anderson suggested that 5.1.2.1 should include a note alerting the reader that the page size field (if nonzero) is binding upon the target even when no page table is used, with a reference to section 9.2 where this requirement is made.

Anderson and Green both suggested that section 5.1.4.5 would be easier to read if the general purpose of the allocate bit was mentioned before going into detail: the allocate bit determines if the ORB is allocating or releasing a node handle. Johansson agreed to rearrange the paragraphs accordingly. Fuller observed that "one or all" would be more clear than "one or more" where the release of node handles was mentioned.

Anderson suggested the text in 7.7.9 could be misread to require the Initiator to take some action on the bus to "terminate the management ORB". Johansson edited the text accordingly. Johansson observed that no guard band was specified for reuse of a timed-out management ORB pointer, to prevent race conditions, and noted that 1394.1 bridges might open a large window for such problems. Johansson suggested that expressing the management ORB timeout

value in ROM gave insufficient flexibility for 1394.1 environments. Anderson agreed with Johansson's suggestion that SPLIT_TIMEOUT would be an adequate guard band on a local bus, but not in a bridged environment. The group deferred discussion to item [7.2].

[6] The group found that topics discussed in recent reflector traffic did not require any changes to the draft.

[7.1] Botchek started a discussion of isochronous services in SBP-3. Anderson described recent products geared towards the video capture and editing market, in which drives or controllers were modified to create DOS filesystems full of DV content on the fly as a connected DV camera is used to record content. Anderson noted that such disks can be connected directly to a computer as a source of media for video editing. Botchek commented that these products were vendor-specific and protocol-specific. Anderson observed that SBP-3 had not yet completed a generic isochronous service model that could have been used in such products, but added that it was unclear if vendors would have used a generic solution even if one had been available.

The group held a brainstorming session and produced the following list of candidate problems to solve in isochronous disks, based largely on past discussions:

1. Content-protected data must remain protected, even if multiple access models are offered.
2. Disk vendors should be able to ship a single product to multiple markets, where the user (or incorporating product) configures the disk appropriately for their desired use.
3. Concurrent access to media - both block access (computer-type filesystem) and isochronous access (real-time media play/record).
4. 1394 TA AV/C Disk functionality model - a disk itself can know how to allocate and find media data and perform real-time record/playback, without hand-holding from a computer.
5. Isochronous play/record function using data in a filesystem managed by a PC - the disk accesses blocks selected by a host PC and uses them for real-time record/playback.
6. Ability to transfer access rights for extents between AV/C disk model and PC model - so both can use direct (native) access for the greatest

performance/features/etc., yet the raw data can be accessed in both ways for maximum functionality.

7. Discovery of extents, such as a PC discovering AV/C extents. This is useful even if the discoverer isn't able to sensibly use all the extents; the discoverer may simply want to convey a list to the user so the user can understand why the disk is full.

8. Avoid the need to draw a fixed partition between PC and AV/C portions of media, as a fixed partition would confine the use of such a product, or reduce the total functionality. (this is perhaps a duplicate of item 2)

9. Permit device performance and capability at least equal to plain ATA/RBC disks, such that a disk meeting this specification could be the sole hard drive in a computer - providing a computer filesystem for the operating system and applications, plus the ability to use 1394 isochronous transfer directly between blocks and other 1394 devices - either with blocks allocated by the PC, or with blocks in extents allocated by the disk.

10. Support a flexible model where products can be created that implement only those items in this list that are of interest to the manufacturer. (but #1 is not optional if protected content can be stored)

11. Optional: Media-level definition of extents such that media can be interchanged between devices - this can work in two (possibly different and incompatible) ways: A bridge and mechanism could be separated from their media (e.g. DVD-RAM), or a bridge could be separated from the mechanism and media (e.g. HDD).

12. Format independence - protocols and commands should not be tied to a single format such as DV, but should be extensible to all future formats. Any particular product may support only certain formats, but the specification should not limit what formats vendors can choose to support.

13. Desirable application example: Use a computer to download a movie from the Internet using a low-quality connection (no real time service), then later use a connected TV to watch the movie with real-time playback (directly from the disk, without copying the data).

14. A capable device could support shared access to an extent: The computer in #13 could be filling an extent at the same time a TV is receiving playback from the same extent - there is no guarantee against an underrun, but we should not exclude simultaneous access when both parties are capable of dealing with the consequences.

15. Capability discovery, such as how many simultaneous streams a device can play/record. This may be static or dynamic, and dynamic is likely to offer more possible combinations of service.

16. Preflight capability, or similar ways to determine if a long operation will succeed or not before undertaking it, or to characterize the resource demands of an operation by simulating or test-running it in whole or in part.

Johansson suggested that most of the items above could be naturally partitioned into three documents:

A. Extent management, with mapping of extents to LUNs and access control - but not block allocation strategies - just bookkeeping of extents (names, properties).

B. SBP, with the minimal isochronous model for play/record and bulk data transfer - but not specific to disks; perhaps usable by scanners or other device types. Specifies ORB-directed isochronous transfer.

C. Command set, perhaps just Play and Record, with hooks for identifying (not specifying) data formats.

The group assigned the 16 items above to Johansson's three categories, as follows:

A. Extent management: 1, 2, 6, 7, 8, bits of 12

B. SBP isochronous: part of 3, part of 5

C. Command set: most of 3, part of 5, most of 12

N/A, other, or enabled by doing the items listed in A, B, and C: 4, 9, 10, 11, 13, 14, 15 (perhaps 15 goes into A or B)

Johansson speculated that B and C, though in separate documents, would be inseparable in a product, while A was independent - in other words, sensible product feature sets would be A, B+C, and A+B+C.

[7.2] Johansson led a review of Bridge-aware SBP-3 target operations:

<ftp://ftp.t10.org/t10/document.02/02-069r1.pdf>

Johansson explained the updates to this document based on discussions at the previous meeting.

Fuller observed that the last two sentences of the first paragraph on page 7 were incorrect and could be removed. Johansson and Botchek agreed. On further study the group determined that the text belonged three paragraphs earlier, in order to prevent an update from being performed when no original login exists. Johansson agreed to make this change and to require the update to check if the login ID value specified is owned by the Initiator requesting the update.

Botchek suggested adding a warning that an update operation might change the login ID. Everyone agreed to this.

Anderson observed that an update login could fail due to a password mismatch. Someone suggested adding a warning that the use of an update login could cause a logout and then fail to establish an updated login. Anderson and Johansson agreed that in the case of such a failure, it was best for the original login to be logged out. Botchek pointed out that update login, no matter what checks it did or did not do, really did not open any new security holes, because any agent with access to the Initiator's Link layer already had total access to the device.

Johansson observed that Annex C (Passwords) replicated a lot of text in the SBP-3 sections covered by 02-069, and suggested that Annex C be revised to eliminate opportunities for discrepancy.

Anderson suggested that paragraph two in 6.4 could mention that FAST_START is required if the corresponding key is found in the Configuration ROM, to match the sense of the surrounding text.

Anderson asked if section 8.1 should mention that a second Net Update must restart any running reconnect_hold timer started by a first Net Update. Fuller observed that the Net Update process, in which the orphan bit changes to one, and then back to zero, was not instantaneous, and suggested that SBP should clearly specify what part of Net Update (start, end, other) should cause the reconnect timer to start. Johansson updated the text at this point, but noted that full details were provided later in the document.

Anderson observed that the table in 6.4 was unclear regarding the effect of FAST_START on an active fetch agent, which is equivalent to a DOORBELL. Johansson agreed to add mention of this effect to the table.

Johansson observed an error in the first NOTE in section 8.2.1, and Anderson expressed confusion about the final sentence in that NOTE. Johansson agreed the text could be improved and said he would revise it.

Anderson observed that the use of "function rejected" status in response to inappropriate bridge-aware (or unaware) login requests might needlessly abstract the nature of the error, making it difficult for an implementer to find and correct the mistake. However, a review of SBP-3r2 found that the "function rejected" status was not used for any other failure condition. Fuller suggested that a single sbp_status value should not be used to report both operational errors such as a resource shortage and also programmatic errors such as an inappropriate login request. Johansson agreed to conduct a survey to see if sbp_status values were used consistently in this regard.

Johansson noted that the two items at the bottom of page 6 should be performed after other items in the list, contrary to the introductory remark claiming the list could be done in any order. Anderson agreed, noting that the reported reason for failing a login should be deterministic even if the login attempt was flawed in multiple ways. Fuller observed that Targets should be careful not to use ack_complete for management ORB requests that might need to indicate an error using some other acknowledge or response value. Johansson said he would add an ordering requirement using the existing order.

Anderson observed that 8.3.1 did not specify how the Target would indicate that it does not support a request for a node handle allocation for a node other than the Initiator. Green observed that the third paragraph on page 9 indicated "resources unavailable", but this error would be an ambiguous way to indicate a lack of support. Johansson agreed to add a general "request type not supported" response suitable for this situation, and to mention it prior to the third paragraph on page 9.

Fuller asked if the "release all node handles" request would release the Initiator's node handle. Anderson observed that 8.3.2 requires all node handles to be released. Fuller and Johansson agreed that the Initiator's node handle should not be released by this request.

Green observed that the login descriptor described in 8.1 should include mention of the Initiator's node handle.

Johansson observed that no requirement was stated to prevent the use of allocate node handle to deallocate the Initiator's node handle. Green said no reason existed to attempt such a deallocation. Fuller suggested that the Target should be required to detect and reject such a deallocation attempt.

Anderson observed that the node handle allocation request can be used to update a previous node handle in the case where the Initiator knows that the referenced node was temporarily absent from the bus. Anderson suggested adding a note that an Initiator using third-party node handles could clean up a temporary disconnection after the ensuing reconnect by using the allocate node handle operation, so the Target would not be burdened with resolving such a situation on its own.

Fuller suggested adding a footnote in section 8.4 to note that the task set is empty immediately after a login, causing the heartbeat timer to start. Anderson noted that a reconnect, abort task set, or (sometimes) abort task would have the same consequence.

Fuller observed that the start of a Net Update should freeze any running timers, and the conclusion of the Net Update should then restart those timers from their initial values.

Action: Johansson to update draft to express this concept where appropriate.

Anderson requested clear text to ensure that a race condition could not cause the heartbeat timer to start running even when the task set was not empty, leading to a surprise reconnect process.

Fuller observed that a write to HEARTBEAT should return a clear error indication if the login has entered the reconnect process or has otherwise become inaccessible. Further discussion led to a change in section 6.4 to encourage the use of `ack_type_error` or `resp_type_error` to indicate any incorrect attempt to access any fetch agent register.

Fuller observed that on a single bus with no bridges but with a bridge aware login, the quarantine bit will always be set. Fuller posited that this situation was not fully addressed in the proposed new text.

Action: Johansson to study required editorial changes.

Fuller commented that the paragraph at the end of page 13 did not fully address the scenario of a Net Update. Johansson agreed to make corresponding edits.

Johansson said he would publish an update (revision 2) of 02-069, and then roll the contents into the SBP-3 draft.

[9] The next scheduled meeting is July 16-17 in Colorado Springs, with T10. The group agreed that the schedule for any additional meetings would be discussed by email or in Colorado.

[10] Johansson briefly reviewed the newly assigned action items.

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>
- 01-069 Steve Powers
Surprise Removal of 1394 Storage Devices
<ftp://ftp.t10.org/t10/document.01/01-069r0.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>

- 01-222 Peter Johansson
Simplified Isochronous
<ftp://ftp.t10.org/t10/document.01/01-222r0.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>

- 01-304 John Fuller
SBP3 Changes
<ftp://ftp.t10.org/t10/document.01/01-304r0.pdf>

- 01-318 Rob Elliott
Elimination of SCSI-2 from SAM-2 SPC-3
<ftp://ftp.t10.org/t10/document.01/01-318r0.pdf>

- 01-330 Peter Johansson
Minutes of SBP-3 Working Group October 3-4, 2001
<ftp://ftp.t10.org/t10/document.01/01-330r0.pdf>

- 01-331 Eric Anderson
Minutes of SBP-3 Working Group November 6-7, 2001
<ftp://ftp.t10.org/t10/document.01/01-331r0.pdf>
- 01-332 Scott Smyers
Isochronous SBP-3
<ftp://ftp.t10.org/t10/document.01/01-332r0.pdf>
- 02-069 Peter Johansson
Bridge-aware SBP-3 target operations
<ftp://ftp.t10.org/t10/document.02/02-069r1.pdf>
- 02-075 Peter Johansson
EUI-48 software interface ID VPD page
<ftp://ftp.t10.org/t10/document.02/02-075r1.pdf>
- 02-206 Eric Anderson
Minutes of SBP-3 Working Group January 21-22, 2002
<ftp://ftp.t10.org/t10/document.02/02-206r0.pdf>
- 02-207 Eric Anderson
Minutes of SBP-3 Working Group March 12-13, 2002
<ftp://ftp.t10.org/t10/document.02/02-207r0.pdf>
- 02-208 Eric Anderson
Minutes of SBP-3 Working Group May 29-30, 2002
<ftp://ftp.t10.org/t10/document.02/02-208r0.pdf>

Latest draft SBP-3 document:

<ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01g.pdf> (change bars/text)
<ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r02.pdf> (no change bars/text)

[end]