DocNum: T10/02-206r0
Author: Eric Anderson
Title: Minutes of SBP-3 Working Group
ftp://ftp.t10.org/t10/document.02/02-206r0.pdf

Minutes of the SBP-3 Working Group meeting, January 21-22, 2002
Outrigger Waikoloa Resort, Hawaii

Attendees:

Eric Anderson	Apple	ewa@apple.com
Lee Farrell	Canon	lfarrell@cissc.canon.com
John Fuller	Sony	jfuller@computer.org
Andy Green	Oxford Semiconductor	andy.green@oxsemi.com
Kashif Hasan	Microsoft	khasan@microsoft.com
Peter Johansson	Congruent Software	Pjohansson@ACM.org
Walt Jones	Microsoft	waltj@microsoft.com
Fritz Nordby	Odd Job Consulting	fritz@2n-1.com
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 Request AV/C expert to define track metadata [Fuller]
4.2 Operational description of login (bus reset) [Johansson]
4.3 Track ID for AV/C disks [Fuller]

4.4 T10 Technical Report for AV Direct-access (AVD)
5. Old business
5.1 Changes to Annex D for AVC encapsulation
5.2 Additional requirements for SBP-3 devices [Fuller]
6. New business
6.1 T10 Technical Report for AV Direct-access (AVD) [Johansson]
6.2 PIMA/PTP Wrapper spec [Anderson]
6.3 Review of reflector traffic [Johansson]
7. Meeting schedule
8. Review of action items
9. 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] The minutes from October 3-4 (Portsmouth) were approved without objection:

ftp://ftp.t10.org/t10/document.01/01-330r0.pdf

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

[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 was active by email, and planned to next meet in January in San Diego.

[3.2] Johansson noted that activity in IEEE 1394.3 has been light, but the standard will soon be ready for a recirculation ballot.

[4.1] Fuller reported no news regarding an AV/C expert to define track metadata.

[4.2] Johansson reported no news regarding an operational description of login (bus reset).

[4.3] Fuller explained that an AVC Object ID would probably be a good way for SBP-3 to reference an object such as a track. The question of uniqueness of Object IDs was still open. Nordby noted that AVC Object ID values must be unique within their list. Nordby also noted that List IDs are unique within a Unit or perhaps a Subunit, so concatenating a List ID and an Object ID would provide a guaranteed unique ID, at least within the subunit. Nordby observed that the largest known List ID is 16 bits, but 32 bits may be used in the future. Fuller said that Object ID is presently specified as 32 bits maximum. Nordby noted that a modern hard drive could hold well over 64K still images, so 32 bit IDs were likely to become necessary soon.

Nordby studied if it is possible to have more that one list of Object IDs on a single volume, and determined that the answer was yes, because the root contents list can have child contents lists, though it's unclear if any AVC devices today make use of this ability.

[4.4] Johansson reported that he created a project proposal for AV Direct-access (AVD) and presented it at the T10 Plenary after the last SBP-3 meeting. Johansson added that document 01-337r0 would be discussed as New Business.

[5.1] Anderson read the notes from the previous meeting. Johansson led a review of Annex D (formerly known as Annex H) to see if the changes made since the previous meeting were correct. Someone observed that figure D-2 needs to add a length field for the size of data written to the response buffer.

Anderson noted that SBP might only allow storing status once per ORB, and said that though it was fine with him to store twice for AVC (to cover Interim) that this

might conflict with other text. Johansson noted that section 5.3 stipulated that status be stored only once per ORB. Johansson showed the encodings for src in status and noted that src 3 would need to be updated. Anderson suggested adding text to clarify that src 0 or 1 status may be stored only once per ORB, and src 3 status may be stored up to once per ORB.

Fuller suggested adding a bit in the transport to suppress interim status when the Initiator has no interest in receiving it. Nordby observed that an Interim can be followed by Rejected, even though an action was performed, so suppression of Interim status might be unwise. The group elected not to add the interim status suppression bit.

[5.2] Anderson read notes from the previous meeting regarding Additional requirements for SBP-3 devices, then presented by Fuller.

Johansson led a review of changes proposed in Monterey and agreed to incorporate them into the draft.

[6.1] Johansson reported that he had proposed a new document 01-337r0 at the most recent Plenary to introduce commands for management of direct-access devices that support streaming media. Johansson then presented a review of the document and asked for feedback. Johansson noted that due to resistance from other T10 members, it would take a focused effort to move forward with the AVD document. Johansson suggested that T10 might be a difficult forum in which to pursue the work, but noted that past history showed that pursuing SCSI commands outside of T10 was also problematic. Johansson asked the group if there was a strong will to proceed with the proposal.

Anderson asked if SBP-3 might face difficulty in approval by T10 if it has no command set to make use of the isochronous features. Johansson said yes, if no command set existed, the isochronous features might not survive in SBP3. Johansson suggested that it would be best to specify at least one command set in order to validate the model. Anderson asked if AVC command sets could work with SBP-3's isochronous services. Fuller said that existing AVC commands worked by themselves, without needing the specific isochronous services of SBP-3.

Johansson noted that the updated READ CAPACITY should be incorporated in the table. Johansson added the topic to New Business and noted that the document should be taken to T10.

Nordby noted that the AVC Disk Subunit states that List ID and Object ID field sizes are described by the implementation and are restricted to 2, 4, or 8 bytes, so the concatenation of List and Object IDs could be 16 bytes. Nordby also

noted that the Root contents list can have Child contents lists, so a single device, subunit, etc., can have multiple lists.

[6.2] Anderson gave a brief review of his plans to specify a standard way to transport the PIMA/PTP (Picture Transport Protocol) over SBP.

Anderson noted that SBP's status block is not large enough for all of PTP's response parameters, but observed that PTP never actually uses all of its defined response fields. Anderson explained that he would provide room for 3 response fields, which satisfied all existing PTP response formats, and that SBP's status block could be expanded if PTP's largest response block was ever expanded. Someone observed that SCSI was defining status messages too large to fit in SBP's status block, so SBP might need to expand this regardless of how PTP is used.

Anderson proposed to store the PTP Device Information outside of the 1394 Configuration ROM, so that this information could be changed without causing a 1394 bus reset. Anderson suggested a quadlet generation count for the data, where interested parties can check for updates.

Johansson reviewed recent email from the SBP-3 reflector.

Green discussed an issue he had raised on the reflector involving the response to multiple Management Agent writes. The group agreed that sending resp_conflict_error as per SBP was best, even if some operating systems would need to be updated to properly handle such responses.

Johansson led a review of 01-180r1. Looking at table 3b (Extent Descriptor), Johansson suggested that Data Format field (2 bytes) should perhaps be larger to allow the inclusion of an IEEE OUI, which would be used like a Spec_ID (SBP).

The group debated whether a computer operating system would want to access the extent directory by reading the entire contents and then looking at each data format, or would want to request a list of extents for each of one or more known data formats.

Johansson suggested that DataFormat could be a protocol identifier or Unit identifier. Anderson said that Unit identifier was fine, but noted that this would tie the spec strongly to 1394, which might be a sticking point for others.

Johansson said he would try to merge QUERY EXTENT into EXTENT MANAGEMENT to reduce the number of opcodes needed.

Anderson asked how the extent management commands would be sent to a disk. Johansson said this was an open question. Fuller suggested LUN 1 be the access point for extent management commands. Johansson suggested that a Target should not allow access to an extent if an exclusive login to another LUN is already established on that extent, and Hasan and Anderson agreed.

Anderson asked if anyone remembered why the ability to change extent sizes was removed. All present agreed that computer filesystems generally could not deal with dynamically sized extents, but could be enhanced to do so. Anderson suggested adding a Set Extent Size selector to the Extent Management command, with suitable warnings about its side effects.

Johansson noted that the Extent Management command 2 needed clarification, as RBC commands were not necessarily the way to access this extent - access should be limited to the LUN through which the command was issued, using the Extent Read/Write commands.

Johansson asked why we had a default RBC extent, noting that modern operating systems would mount filesystems on all RBC extents discovered. Anderson said that it had been desired that LUN 0 always correspond to an RBC LUN for legacy support and noted that the Set Default RBC Extent command allows LUN 0 to be directed to other extents if desired. Anderson observed that Extents did not generally correspond to LUNs. Someone asked what would happen in the default RBC extent was changed while someone was logged into LUN 0. Fuller observed that RBC Extents have to map to LUNs in order to be accessed, while other kinds of Extents (eg AVC) did not have to correspond to LUNs.

[7] The April meeting was canceled due to a conflict with the 1394 Trade Association meeting in Barcelona.

Johansson asked if one RBC extent would be sufficient for any device. Anderson said that as long as the architecture didn't make it difficult to support multiple RBC extents in the future, one seemed adequate for the present, and Hasan agreed.

Johansson asked if the extent management functions should be accessed through a dedicated LUN (perhaps LUN 1) or through management ORBs sent to other LUNs. Johansson noted that management ORBs are inconvenient because they are single-threaded.

Johansson noted that ORB-driven access to extents would have to go through some LUN, and that individual LUNs per open extent would be needed if exclusive access to extents was to be supported. Johansson asked if these

LUNs should be created on demand, and be ephemeral. Anderson and Fuller agreed that Lock/Unlock should be persistent, but LUNs could be ephemeral.

Johansson noted that with dynamic LUNs, the Read/Write Extent relative commands were no longer needed, because the LUN created for an extent would already be relative to that extent. Anderson noted that dynamic LUNs might be left out of the 1394 Configuration ROM, to avoid requiring a bus reset, and to avoid accidental discovery by other initiators between their creation and the subsequent login. Fuller suggested the login be automatic, eliminating the need for a LUN, but Anderson and Johansson noted that existing software would be difficult to adapt to such use, and it would complicate the internal workings of a device. Johansson suggested that dynamic LUNs could be persistent since they were not visible in the Configuration ROM, but Anderson noted that initiators would have to "fish" for a previously known LUN, and would not be able to tell if the LUN they found was really the same one they had previously created.

Johansson suggested that Create Extent can be used to change extent sizes (perhaps after renaming it) and added that the parameters would be size, and (for an existing extent) extent ID. Someone asked if split/join should be supported for extents. Anderson said that if supported, split/join should be allowed only for native extents. Non-native extents might be able to be split or joined through other command sets, such as AVC.

Johansson noted that the modified read capacity was no longer needed, because the extent data set would give the size for each extent, and ordinary read capacity could be used within an RBC extent. Johansson suggested having an extra entry in the extent data set that corresponded to the entire disk, so that a new command to learn the size would not be necessary. Nordby noted that another synthetic extent could indicate the free space on the disk, with no other commands allowed on the free extent. Johansson noted that the EXTENT DIRECTORY data structure could be arbitrarily expanded to hold the total and free size information, which would be more concise and perhaps less misunderstandable than using synthetic extents.

Johansson proposed a map operation in which one specifies an extent ID, and gets a LUN in response. Johansson asked if the map should specify the access control. Anderson noted that the subsequent login would specify exclusive or not. Green suggested combining the lock/unlock functions with the map function. Fuller and Anderson pointed out that the map/lock operations are related but not the same, and there might be uses for them independently.

The proposed parameters for the map command were extent ID and a read/write/readwrite selector. A mapping would be unique - one cannot map one extent to two LUNs. If a mapping already exists, the existing LUN is returned.

Fuller suggested that unmap is not needed, because mappings can expire if some time elapses during which nobody is logged in (or eligible to reconnect). This time could be the reconnect time, even though a reconnect is not actually happening. Nordby suggested that mappings should not expire as long as the Initiator is logged in to the master LUN (LUN 1).

Fuller noted the problem of Initiators who want to take turns using a LUN; presently there is no way to get notified when someone releases a login. Fuller suggested that the same problem might apply to extents.

Anderson said that the master LUN (1) should be at least encouraged to support multiple logins, and logins to the master LUN should be discouraged (or perhaps disallowed) from being exclusive, and Jones agreed.

Johansson suggested that mappings should be persistent across power cycles. Fuller said that persistence would incur extra cost, or delay for accessing drive media. Anderson noted that an Initiator needs the ability to discover and map an extent at least once, and could always repeat this step rather than relying on persistent LUNs. Jones agreed with Anderson, and agreed with Fuller that the important issue was to consider the burden upon the Target. Green said that as a Target implementer he felt it would be easier to automatically dispose of unused mappings (perhaps on a lazy basis) than to maintain persistent mappings.

Johansson noted that without persistence, it was possible that an initiator could create a mapping, get distracted, and then perform a login to the LUN - only to find that a different extent was then mapped to that LUN. Anderson suggested a tedious verification process that initiators could use. Anderson then suggested that mappings must stay valid until they are first logged in to, or a power cycle. Anderson modified this suggestion to create a reference count - performing a map operation increases a reference count; logging into the LUN decreases the reference count. This way, an Initiator that desires to repeatedly log into the same LUN can perform the map to make the mapping sticky and prevent it from expiring. Jones noted that a map could be performed immediately after a login to ensure that a future login would be possible. Johansson suggested a verification step to ensure a login had accessed the desired LUN, as an alternative to reference counts. Anderson noted that his tedious process did that, by just issuing the map command again to verify that the same LUN was returned, and that without the time measurement part it would be a simple technique. Anderson noted though that the technique was optional, so Initiators might be likely to skip it, leading to reduced reliability.

Fuller said that a leak of maps, even if rare, was undesirable if only a power cycle could clear it, because some implementations will never encounter power cycles.

Johansson noted that targets already have timers (for reconnect) so they can be required to time other things too. Johansson suggested that idle mappings expire after a generous amount of time. Fuller noted that targets need not be required to actually expire such an idle mapping until they need to reclaim it to satisfy a subsequent mapping request.

Anderson noted that with this scheme, no reference counts or verification would be needed.

Jones noted that Query Extent could be made to return information about existing mappings, so that a second map operation wouldn't be needed (it would restart the timer).

The group discussed the issue of how to set channel numbers and associate I/O with an ISO/IEC 61883 Plug register and the corresponding Connection Management Protocol, when those protocols were used. It was observed that the plugs could exist on the disk (Target) node, in a fully autonomous device, or could exist on the Initiator, with the Target simply providing the data movement, but the Initiator owning management responsibility for the transfer. Johansson said he might write a new proposal to replace 01-180r1.

The group determined that Record and Play commands might need to set or learn the following: 1394 isochronous channel numbers, 61883 Plug numbers, data format information for timestamp adjustment, timing information such as a cycle number or sync bit value on which to start transfer, and some kind of play list to define the actual data to be transferred.

Green asked about multichannel record and/or playback. Anderson noted that Smyers had originally proposed such an ability, but it seemed to have been lost is subsequent discussion. Johansson said that information about the Create Stream command is strewn over various drafts but isn't coherent at the moment.

Johansson said he would update the draft with the src=3 text from 01-287r1 and related text.

Green suggested moving the June 5-6 meeting to May 29-30, and nobody objected.

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

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

01-332 Scott Smyers
Isochronous SBP-3
ftp://ftp.t10.org/t10/document.01/01-332r0.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

Latest draft SBP-3 document:

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

[end]