Minutes for RAID working group 03/15/94
Accredited Standards Committee
X3, Information Processing Systems

Doc. No.:
X3T10/94-062R0
Date: March 15, 1994
Project:
Ref. Doc.:
Reply to: G. Penokie

To: Membership of X3T10

From: George Penokie and Ralph Weber

Subject: Minutes of RAID Study Group Meeting
March 15, 1994

Agenda

1. Opening Remarks
2. Attendance and Membership
3. Approval of Agenda
4. Report on last RAID Working Group
5. Dual-Controller Issues
6. SCSI Disk Array Model (94-040r2)
7. SCSI-3 Addressing (94-031r2)
8. SDA Commands and Mode Pages (94-042r2)
9. SDA States and Types (94-041r2)
10. Error Handling for SCSI Controllers (94-024r2)
11. Action Items
12. Meeting Schedule
13. Adjournment

Results of Meeting

1. Opening Remarks

George Penokie the RAID Study Group Chair, called the meeting to order at 9:08
am, Tuesday, 15 March 1994. He thanked Qlogic for hosting the meeting.

As is customary, the people attending introduced themselves. A copy of the attendance list was circulated for attendance and corrections.

It was stated that the meeting had been authorized by X3T10 and would be conducted under the X3 rules. Ad hoc meetings take no final actions, but prepare recommendations for approval by the X3T10 task group. The voting rules for the meeting are those of the parent committee, X3T10. These rules are: one vote per company; and any participating company member may vote.

The minutes of this meeting will be posted to the SCSI BBS and the SCSI Reflector and will be included in the next committee mailing.

George stated that this is the 17th meeting of the RAID study group and the 4th joint RAID Advisory Board/RAID study group meeting. The purpose of the group is to deal with interface issues related to using RAIDs. The study group will assess the issues and then formulate a strategy for dealing with them.

2. Attendance and Membership

Attendance at working group meetings does not count toward minimum attendance requirements for X3T10 membership. Working group meetings are open to any person or company to attend and to express their opinion on the subjects being discussed.

The following people attended the meeting:

RAID Study Group Meeting Attendees

Name	S	Organization	Phone Number -or- Electronic Mail Address
Mr. Randy Hall		Array Technology	303-938-6035
Mr. Thanh Pham		AST Research	714-727-6884
Mr. Charles Binford	A	AT&T/GIS	316-636-8566
Mr. M.W. Jibbe	A	AT&T/GIS	316-636-8810
Mr. Jeff Brown		Clariion-DG	508-898-5006
Mr. Ed Haske		CMD Technology	714-454-0800
Mr. Allan Cole		CORE International	407-997-6044 x505
Mr. Bill Dallas	A	Digital Equipment Corp.	dallas@wasted.enet.dec.com
Mr. Ralph Weber	A	Digital Equipment Corp.	weber@star.enet.dec.com
Mr. Mitch Rosich		EMC Corp.	rosich@emc.com
Mr. Howard Grill		Formation	609-234-5020
Mr. Bill Hutchison		Hewlett Packard Co.	hutch@.boi.hp.com
Mr. Andrew Hill		HI-Data	617-937-6770
Mr. Paul Boulay		Hitachi Computer Products	408-986-9770
Mr. Ken Cummings	A	IBM Corp.	602-799-2182
Mr. Giles Frazier	A	IBM Corp.	512-838-1802
Mr. George Penokie	P	IBM Corp.	gop@rchvmp3.vnet.ibm.com
Mr. Kevin Pokorney		Intellistor/Fujitsu CPA	303-682-6649
Mr. Dave Fellinger		Mega Drive Systems	310-247-0006
Mr. Dick McCormick		Peer Protocols	714-476-1016
Mr. Herb Silverman		Peer Protocols	714-476-1016
Mr. Tony Fung		Storage Concepts	714-852-8511
Mr. Frank Ruthofford		Storage Concepts	714-852-8511
Mr. Atul Talati		Storage Concepts	714-852-8511
Mr. Robert N. Snively	P	Sun Microsystems, Inc	bob.snively@eng.sun.com
Mr. Joe Molina		Technology Forums	507-931-0967
Mr. Gary Watson		Trimm Industries	trimm@netcom.com
Mr. Peter Dougherty		Unisys	714-380-6270
Mr. Duck	D	Hyatt pool	
Mrs. Duck	D	Hyatt pool	

28 people present
2 ducks present

Status Key: P - Principal ** X3T10 status still being defined
A - Alternate
O - Observer
L - Liaison
S,V - Visitor
D - Duck

3. Approval of Agenda

The agenda developed at the meeting was approved.

4. Report on last RAID Working Group

Last RAID Working Group was held in Austin, TX. Several members of the RAID Advisory Board were present. George stated his intention that the minutes of the Austin meeting be published on the SCSI Reflector. Questions were raised about whether that actually occurred. George promised to review his minutes publication practices.

5. Dual-Controller Issues

George noted that the C-LUI concept may resolve most or all of the dual-controller issues raised at previous meetings. The C-LUI concept is discussed in the new SCSI-3 Disk Array Model document (see below). George reported that Giles Frazier (IBM, Austin) and Doug Hagerman (Digital Equipment Corp.) both have reviewed the SCSI-3 Disk Array Model from the perspective of dual controller issues. Giles has some small additions that should be added. Doug believes that C-LUIs cover all his concerns. George is waiting for the dust to settle.

Giles and George discussed this issues that Giles has raised. The issues revolve around the master/slave aspects of dual controller implementations. How does a master exchange its ownership control with a slave? What commands are used? How do addressing issues (LUN 0) get managed?

At this point, the discussion turned to a picture of a specific single controller configuration. The various P-LUI and C-LUI components were identified, described (as to purpose), and named. Then, the single controller configuration was mirrored into a dual controller configuration. A result of this was broadening George's experience in dual controller configurations.

By the end of the discussion, George was convinced that most of the dual controller issues are addressed by the current C-LUI concepts. However, George felt that additional descriptive text is necessary in the Disk Array Model to clarify C-LUI application to dual controller configurations.

6. SCSI Disk Array Model (94-040r2)

George presented a revised version of the SCSI-3 Disk Array Model (X3T10/94-040r2).

The major new concept in the revision of the SCSI-3 Disk Array Model is C-LUI. A C-LUI is a Component Logical Unit Identifier. A C-LUI represents an entity that is not identifiable as a SCSI-3 Device Type. To properly contrast with this, the definition of P-LUI has been changed to make P-LUI represent an entity that is identifiable as a SCSI-3 Device Type.

George walked through the Model discussing the specific effects of the C-LUI addition. At several points, corrections were noted to improve the details of C-LUI definition and usage.

George noted a long-standing convention that he and Doug Hagerman have

discussed. The Model describes redundancy groups as being attached to spares.
Doug prefers attaching spares to redundancy groups. Recently, George realized
the conceptual benefits of using Doug's attachment model. However, George believes that the benefits of long-term consistence in the Model outweigh the conceptual simplification provided by reversing what is attached to what.

Questions were raised regarding relationships when A is attached to B and B is
attached to C. Can A be used as a path to gain information about C? C-LUIs complicate this problem, because C-LUIs can be attached to other C-LUIs. Using C-LUIs as a mechanism for describing dual controller configurations also
may complicate this issue. George was not certain that a problem really exists. George promised to work on the issue when/if its status becomes clearer.

As discussion progressed, the possible need for a new service was uncovered. The new service would be directed to a P-LUI. The service could be named "Report P-LUI Attachments". The service would ask the P-LUI to report everything that is attached to it. Upon further reflection, an equivalent need was discovered for C-LUIs. George agreed to add the service to the next revision of the Model.

Questions were raised about what parts of the standard will be optional and mandatory. George felt that exact definition of optional and mandatory should
be left as one of the last steps in the standardization process. If pressed to guess, George felt that all the commands and services would be optional. The addressing and states definitions might be mandatory.

George presented a list of possible C-LUI devices (things). The list was extracted from a private communication that George received from Doug Hagerman. George encouraged those attending to send him their lists.

7. SCSI-3 Addressing (94-031r2) Penokie

George presented a revised version of the SCSI-3 Logical Unit Addressing proposal (X3T10/94-031r2). Concerns were raised about the relationship of this P-LUI addressing scheme and the vendor-specific schemes already in use. George noted that the existing vendor-specific schemes will be valid forever, but he assumes that vendors will gradually convert to the standard (generalized) scheme.

Other concerns were raised about the reduction of addressable volume set LUNs from 64 to 32. George noted that the reduction has value, in the form of the ability to address V-LUIs and P-LUIs without doing address remaps.

The level of change in 94-031r2 is very small. George declared the document stable. He said that it will be presented to the X3T10 Plenary for approval as soon as the SCC project proposal is approved by X3.

8. SDA Commands and Mode Pages (94-042r2)

George presented a revision to the document defining the new commands and mode
pages needed to implement SCSI disk arrays (X3T10/94-042r2). The major source

of changes was the addition of the C-LUI concept.

9. SDA States and Types (94-041r2) Penokie

George presented the latest revision of the document defining states of an SDA and the types of devices needed within Disk Arrays (X3T10/94-041r2). There was some discussion of the "base address" wording as opposed to the "P-LUI 0" wording. In fairly short order, George was convinced to adopt the "base address" terminology.

George noted several changes resulting from the new C-LUI concept. The entire list of C-LUI states is new. George solicited proposals for C-LUI states in addition to those that he has listed.

10. Error Handling for SCSI Controllers (94-024r2) Hagerman (Weber)

Substituting for Doug Hagerman, Ralph Weber presented the latest draft of the Error Handling for SCSI Controllers document (X3T10/94-024r2). The document was new to most people present. Ralph describe how the lists of new ASC/ASCQs on pages 7 through 9 were taken from the vendor-specific ASC/ASCQs of several products that Digital markets. Ralph also noted the addition of a Disk Array device type column for the existing ASC/ASCQs.

George said that he wants to scan the lists for duplicates. George described the normal policy for defining new ASC/ASCQs, which includes a duplicates check (made by George). No one else had scanned the list for codes of interest to them. Some real-time scanning occurred during the meeting.

Some people present requested definition of some ASC/ASCQs that represent x-LUI states (particularly failure-mode states). George said that he will consider that idea.

It was noted that Jerry Fredin (AT&T/GIS) has a list of ASC/ASCQs based on contributions from manufacturers at the time the RAB was formed. Although the list is somewhat moldy now, it still is deemed to be a valuable reference. Ralph was asked to have Doug contact Jerry.

11. Action Items

a)	Penokie	Revise the SCSI Disk Array Model (94-040) document based on input received and toward making it cover dual controller configurations.
b)	Penokie	Revise the SDA States and Types (94-041) document.
c)	Penokie	Revise the SDA Commands and Mode Pages (94-042) document to account for the "report attachments" service.
d)	Hagerman	Contact Jerry Fredin (AT&T/GIS) for original manufacture inputs regarding vendor-specific ASC/ASCQ definitions for RAID.
e)	Hagerman	Revise the Error Handling for SCSI Controllers (94-024) document.

12. Meeting Schedule

The next meeting of the RAID Study Group is planned for April 19, 1994 at the Hyatt Regency in Bellevue, WA. The meeting is expected to run from 9:00am-5:00pm. This meeting will be a joint RAID Advisory Board Host Interface Group and X3T10 RAID Study Group meeting.

13. Adjournment

The meeting was adjourned at 3:11 pm. on Tuesday, 15 March, 1994.