

Accredited Standards Committee  
X3, Information Processing Systems

Doc. No.: X3T10/94-108r0  
Date: May 17, 1994  
Project:  
Ref. Doc.:  
Reply to: G. Penokie

To: Membership of X3T10  
From: George Penokie and Ralph Weber  
Subject: Minutes of RAID Working Group Meeting - May 17, 1994

## Agenda

1. Opening Remarks
2. Attendance and Membership
3. Approval of Agenda
4. Report on last RAID Working Group
5. Procedural and Process Notes
6. SDA Commands and Mode Pages (94-042r4)
7. SDA States and Types (94-041r4)
8. Dual-Controller Issues
9. SCSI Disk Array Model (94-040r4)
10. Mandatory/Optional RAID Requirements
11. Action Items
12. Meeting Schedule
13. Adjournment

## Results of Meeting

### 1. *Opening Remarks*

George Penokie the RAID Working Group Chair, called the meeting to order at 9:08 am, Tuesday, May 17, 1994. He thanked AMP 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. Working group 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 19<sup>th</sup> meeting of the RAID study group and the 6<sup>th</sup> joint RAID Advisory Board/RAID working 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.

### RAID Study Group Meeting Attendees

Name	S	Organization	Electronic Mail Address
Mr. Larry Lamers		A Adaptec	LJLamers@AOL.com
Mr. Ken Scherzinger		Amphenol Spectra-Strip	
Mr. Randy Hall		Array Technology	hall@arraytech.com
Mr. Jerry Fredin	A	AT&T/GIS	Jerry.Fredin@WichitaKS.NCR.COM
Mr. George Su		CMD Technology	714-454-0800
Mr. Allan Cole		CORE International	73532.1636@compuserve.com
Mr. Bill Dallas	A	Digital Equipment Corp.	dallas@wasted.enet.dec.com
Mr. Doug Hagerman	A	Digital Equipment Corp.	hagerman@starch.enet.dec.com
Mr. Paul Massiglia	O	Digital Equipment Corp	massiglia@genral.enet.dec.com
Mr. Ralph Weber	A	Digital Equipment Corp.	weber@star.enet.dec.com
Mr. David Race		DPT	
Mr. Dal Allan	P	ENDL	250-1752@MCImail.com
Mr. Howard Grill		Formation	609-234-5020
Mr. Bill Hutchison		Hewlett Packard Co.	hutch@.boi.hp.com
Mr. Paul Boulay		Hitachi Computer Products	408-986-9770
Mr. Giles Frazier	A	IBM Corp.	gfrazier@ausvm6.vnet.ibm.com
Mr. Paul Hodges	A	IBM Corp.	phodges@vnet.ibm.com
Mr. George Penokie	P	IBM Corp.	gop@rchvmp3.vnet.ibm.com
Mr. Dick McCormick		Peer Protocols	714-476-1016
Mr. Gerald Houlder	A	Seagate	gerry_houlder@notes.seagate.com
Mr. Gene Milligan	P	Seagate	gene_milligan@notes.seagate.com
Mr. Joe Argento		Seek Systems	
Mr. Vit Novak	A	Sun Microsystems, Inc	vit.novak@eng.sun.com
Mr. Robert N. Snively	P	Sun Microsystems, Inc	bob.snively@eng.sun.com

24 people present

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

## 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 Bellevue, WA. Several members of the RAID Advisory Board were present. George stated that the minutes of the Bellevue, WA meeting have been published on the SCSI Reflector.

## 5. Procedural and Process Notes

George announced that the SCSI-3 Controller Commands (SCC) project proposal has been approved. This means that the X3T10 study group now is a working group. It also means that the working group can take proposals to the X3T10 Plenary meeting for votes. The details of this process were discussed.

The group discussed which RAID documents could be voted on at the Thursday Plenary meeting. George was certain that the addressing document could be voted on, because it has not changed for months. He was less confident that attempting to vote on the model would survive a two-week-rule challenge.

George announced that he has converted to using Frame Maker. There was a brief discussion of how document distribution and interchange will be handled in this environment.

The group discussed the best approach for getting the document through the X3T10 voting process. George is anticipating a vote for a forwarding letter ballot at the September Plenary. Dal suggested an introductory session at either the July or September Plenary. George said that he may not be able to prepare the introduction and do the necessary editing by July. George was more comfortable with doing the introduction in September.

## **6. SDA Commands and Mode Pages (94-042r4)**

George presented a revision to the document defining the new commands and mode pages needed to implement SCSI disk arrays (X3T10/94-042r4). While trying to assign SCSI op codes to the new commands, George decided that too many SCSI commands would be required. So, George grouped the commands into four groups:

MAINTENANCE, REDUNDANCY GROUP, SPARE, and VOLUME SET. Each group was assigned a single SCSI op code. A Service code field is defined in each SCSI CDB giving the full array command set, as previously defined by the group.

Questions were raised about whether different Service codes cause a single SCSI op code to either transfer data in or out. George said that the data transfer direction would depend on the Service code. Bob Snively requested that data transfer direction be defined by the SCSI op code. George agreed to break the op code by data transfer direction, yielding: MAINTENANCE IN, MAINTENANCE OUT, REDUNDANCY GROUP IN, REDUNDANCY GROUP OUT, SPARE IN, SPARE OUT, VOLUME SET IN, and VOLUME SET OUT.

George began discussing the individual op-code/service-code commands. Questions were raised about the extensibility of some fields. George, Doug Hagerman, and several others agreed to consider extensibility after the basics of the command set have been determined.

George described every command and its parameters. Questions of detail were raised for a few commands. George was asked to increase the reserved space in the headers of the parameter data on several commands.

Toward the end of the review, Giles Frazier complained about the usage of command to name a CDB whose function is qualified both by op code and service code. George agreed to consider a name that uses the word "service" instead of "command."

In Table 51, several important concepts were raised. P-extents should be rebuilt from a specific redundancy group or from an SDA selected (any) redundancy group. The specific/any distinction should be selected using a bit. The parameter data should contain only one R-LUI per P-extent. Doug Hagerman reminded the group that the P-extent (in this case) may cover several redundancy groups (because the P-extent is a free-form command parameter). Ultimately, this led to the realization that the original Rebuild P-Extent parameter list format was correct.

As the discussion continued, George put more red ink on the Table 51 slide than on any previous slide. The specific needs of the rebuilding operations involving multiple redundancy groups (that may overlap) produced several conflicting definitions. A major concern was whether the rebuild command would fail when a redundancy group could not be rebuilt.

Dal Allan proposed two versions of the rebuild: 1) rebuild all and 2) rebuild all except. Both of these provide a paradigm that allow the SDA to report an error if the rebuild cannot be completed. Dal noted an additional model; rebuild everything that can be rebuilt using this list of redundancy groups. This most closely matches the command, as originally proposed. However, it can be viewed as incomplete because the SDA cannot report a failure to perform the rebuild. The SDA gets into this trap because it must assume that the host intentionally omitted the redundancy group(s) that prevented the complete rebuild.

The group agreed on the following rebuild operational rules:

- a) all assigned space within the P-extent shall be rebuild, otherwise there is an error
- b) the ways of specifying the redundancy groups to be used are:
  - 1) rebuild with ALL redundancy groups (if overlaps, use ALL)

- 2) rebuild with ANY redundancy group (if overlaps, use ANY one)
- 3) rebuild with all redundancy groups, except those listed (if overlaps, use ALL)

During a discussion of continuous verification of V-LUI ranges, George noted that there is no way to report if continuous verification is enabled. For a few brief moments, George was ready to add reporting to one of the report commands. Then, the complexity of that hit home, and George said that no reporting would be provided.

## **7. SDA States and Types (94-041r4) Penokie**

George presented the 4<sup>th</sup> revision of the SDA States and Types document.

George noted that the Ready bit is an or condition of the states of all logical units on the SDA. George also noted several reporting situations where the exact nature of the reporting methodology is vendor specific.

George was asked to define a command that returns states of all addressable things in an SDA. The command needs to return lists of states. George agreed to develop the details before the next meeting.

In the set of redundancy group states, George agreed to add a new state, Invalid Protected Space. Invalid Protected Space means that the protected space is no longer intact.

George noted that he removed the Exposed state from the set of spare states.

Spares cannot be in the exposed state.

## **8. Dual-Controller Issues**

Jerry Fredin (AT&T/GIS) presented a collection of questions regarding how the SDA model applies to dual-controller configurations. His first question concerned how paths into the controller are addressed. When a controller (or DACL) fails in a redundant configuration, how does that failing unit get identified for maintenance personnel?

The discussion turned to using the get states command (described above). Ralph proposed definition of a new state for controller C-LUIs, I'm talking to you (ITTU). Used in combination with the get states command, the ITTU state allows a host to map relationships between multiple-controller C-LUIs and the addresses that it asserts to get to them. This would further help reconfiguration because it would allow the host to quickly build commands that have the effect of saying, "Provide access to R-LUI five through the C-LUI that I'm talking to."

The group also discussed some kind of "Instruct C-LUI" command. This was felt to be a sub-command the MAINTENANCE commands. Ralph proposed a command format that George will consider.

## **9. SCSI Disk Array Model (94-040r4)**

George reviewed changes to the SCSI Disk Array (SDA) model. The changes were minimal, but George still received some comments about them.

## **10. Mandatory/Optional RAID Requirements**

The group discussed the process by which they will define what commands and/or features are required for SCSI Disk Arrays. The need for very simple implementations (that don't do everything) was discussed. However, there also is a market need for more complex implementations. No resolution was reached on the specific requirements. George expressed a preference for resolving the mandatory/optional issues at a future meeting. Doug pushed for simplicity; ideally all commands mandatory. Dal proposed dividing implementations into configurable and non-

configurable. George agreed that an INQUIRY data bit would be defined to segregate configurable and non-configurable DACLs.

## **11. Action Items**

- a) Penokie Revise the SDA Commands and Mode Pages (94-042) document based on comments and discussion topics from the meeting.
- b) Penokie Define a command for returning states (see notes in the SDA States and Types discussion).
- c) Penokie Revise the SDA States and Types (94-041) document.
- d) Penokie Roll SDA Model, SDA Commands and Mode Pages, and SDA States and Types in to a single document.
- e) Hagerman Propose ASC/ASCQ definitions for RAID. Said proposal must be sent to X3T10 for approval and inclusion in the SPC.

## **12. Meeting Schedule**

The next meeting of the RAID Study Group is planned for June 21, 1994 at the Sofitel in Minneapolis, MN. 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 Working Group meeting.

## **13. Adjournment**

The meeting was adjourned at 5:31 pm. on Tuesday, May 17, 1994