

Doc. No.: X3T10/95-313r0

Date: September 21, 1995

Project:

Ref. Doc.:

Reply to: Steve Sicola

To: Membership of X3T10

From: Steve Sicola, Chair HA Study Group

Subject: Minutes of X3T10 High Availability Study Group  
Bedford, NH -- September 11, 1995

### Agenda

1. Review of HA Preliminary Profile for Fault Tolerant Controller Configurations
2. Discussion of any/all specific changes required(if any) to SCC to make Profile useable.
3. Discussion of Doug Hagerman's High Availability Disk Drive Specification

#### Attendees:

Ms Nancy Cheng  
Mr. Doug Hagerman  
Mr. Charles Monia  
Mr. Robert Snively  
Mr. Ed Quiet  
Mr. Stephen J. Sicola  
Mr. Rod Dekoning

The meeting opened with a change in agenda, while a number of other attendees were still showing up. Doug opened up with an overview of his high availability disk drive specification. It covers high availability aspects such as how disks react to bus resets, errors, hot plugging, etc. The document was distributed to those present and it was decided that the document should be further reviewed at the next HA study group meeting.

The review of the HA Preliminary Profile for Fault Tolerant Controller Configurations began with a few slides from Steve Sicola reviewing the general theory of operation of fault tolerant controller configurations within the confines of SCC. This included the single 'logical' controller model with multipel ports, which could physically be multiple controllers, so long as all controllers shared the same devices. The aspects of setting up attachments and reporting attachements were discussed as well.

Rod Dekoning of Symbios and Bob Snively of Sun brought up points about the wording of sharing the same devices. It is safe to say that the controllers must share the same path to all devices instead of sharing all devices. This opens up a few more possibilities for fault tolerant configurations.

Other items brought up early in the discussion were that the profile must allow for preconfigured systems with attachments already in place and that the profile must support active/standby configurations to round out the profile. The detection mechanisms for controller failure can be host or controller based. The host can certainly still detect a failure by timing out, but the addition of two ASC/ASCQs that describe Failure or Failback(controller returning to the configuration) will allow for higher availability of customer data from controller configurations. Also, the use of LUN0 in the profile should be changed to the 'base address of the controller.'

It was noted early by Steve Sicola that the SCC attachment command had no way of identifying the configuration once the attachment was made as well as the fact that there are no fields in the Attach to Component Device command to identify the WWN's of the controllers involved in the attachment. This led to a long discussion between Bob Snively, Rod Dekoning, and other attendees as to the merits of identifying configurations and the use of SCC to setup and report them. After a long discussion, it was decided by those present that the benefits of registering the configuration, reporting the configuration, and detecting of failures and new members of the configuration would be

of great benefit to generalize the use of fault tolerant controllers within many host SCSI drivers. It was at this point that specific changes to SCC were discussed.

The changes discussed were those around the Attach to Component Device and Report Attachments. Many options were discussed to cover the possible configurations, but the issues of initial setup, repair/replacement, and upgrade of the configuration drove the need for the attach to component device command to include input: who to attach to, type of attachment

output: Name of attachment (suggested to FC naming convention of 8 bits vendor unique, 24 bits user prog)

The uses of the Attach to Component device is in configuration and ease of verification of configuration.

The Report attachments would need the following:

input: Name of Attachment

output: WWN's of all attached components, where the 1<sup>st</sup> WWN reported is the one who received the report attachments command Followed by: WWN's of eligible components that could be added to attachment list.

The uses of Report attachment command would include host crash, host return, new hosts, new controller, replaced controller.

The Detection of failures is basically an issue of host polling vs. controller direct reporting (ASC/ASCQ's). The addition of the two ASC/ASCQ's will allow for both methods to work properly and be allowed. The HA group defers to the PFA rules on exception handling to drive solutions beyond this.

The meeting concluded with the decision that the specific changes to the SCC document should be redlined for the next Working Group meeting in Palm Spring in November. Steve Sicola will write these up and distribute before the meeting for early review and for any possible modifications to the redlines for completeness and accuracy.