

Stewart Wyatt, Secretary HP 208-396-3594, stewart_wyatt@hp.com

1. Introductions: Group

Facilitator Dale LaFollette called the meeting to order shortly after 1 PM. Dale noted the name change for the group as is reflected in the heading of these minutes. He had the members of the group introduce themselves.

2. Approval of this agenda, T11/99-523v1: Group

Approved

3. Approval of 8/4/99 minutes, T11/99-471v0: Stewart Wyatt

Approved

4. Review of old action items: Stewart Wyatt

Old forwarded action items

- #1. Eric Oetting - to collect issues from the SCC letter ballot comments for the SPC-2. - Completed
- #2. Stephen Gold (Stewart Wyatt) HP: Resolve IBM letter ballot comments on TapeAlert. - Completed
- #3. Bob Snively and Bob Kembel: Create a list of the specific BLS and ELS which are allowed before Login including explicit Login in an annex for later placement in the FC-FS. - Ongoing
- #4. Bob Snively: Review the SRR statement, it may not be flexible enough for the target. SRR must be properly interpreted in context and command known by the target. "Nasty case is did the transfer ready or the or response (to an error) get hit." - Completed
- #5. Bob Snively: Return clause 2.8 to previous text. - Completed
- #6. Bob Snively: Change Table 2 - Discovery of FCP capabilities entry "Initiator performs REC", Discovery mechanism needs to be changed from "None/Process Login" to "See References". - Completed
- #7. Dave Peterson and group: Extend current document to cover LIP discovery process. Participants asked to communicate concerns to Dave. - Completed
- #8. Carl Zeitler: Review need for process associators. - Due in October.

New Action Items as of last month

- #1. HP (Steve Jerman/Sid Crighton): Continue MAM discussion. Post presentation made today. Contact Ralph Weber about including MAM in the SPC-2. Make presentation to SCSI working group next month. Resolve issues with Sony AIT per Joe Breher's concerns. Determine need, if any, for content to be placed in the SSC. - Completed
- #2. Bob Snively: (Binford FCP comment #006) LUN reset should be added to the table 3.
- #3. Bob Snively: (Binford FCP comment #007) FCP actions should only affect FCP exchanges.
- #4. Bob Snively: (Binford FCP comment #014 and #015) Check to see that Bus Inactivity limit and Disconnect time limit are specified consistently with other SCSI standards.
- #5. Bob Snively: (Binford FCP comment #022) Define what a target shall do when it receives a ABTS.
- #6. Bob Snively: (Binford FCP comment #027) Check wording to see that REC data is counted is over an exchange not a sequence.
- #7. Bob Snively: (Binford FCP comment #029) Optional error recovery procedures need to be clearly marked. The requirements for using error recovery procedures in class 3 need to be clearly specified for disk drives. A target may request an FCP_CONF at its disgression.
- #8. Bob Snively: (Binford FCP comment #033) Target must use valid RX_ID (<> FFFFh) when using FCP_CONF.

#9. Bob Snively: (Binford FCP comment #018) Add a default value for RR_TOV for non REC_TOV error recovery.

#10. Bob Snively: Binford FCP comments #008, using WWN for reservations, and #010, title for Confirmed Completion Allowed, appeared to be unresolved without specific action items. In addition the proposal by Jim Coomes to clarify the DSA field appeared unresolved as well as a proposal to rename it more appropriately.

Status of items #2 to #10, Bob believes these issues have all been addressed and will be resolved in the next revision of the FCP.

5. SSC, T10 Working Drafts SSC-R19, LB Comment Resolution T10/99-228r1: Dave Peterson

Resolution of letter ballot comments

IBM Comment 96. Clause 5.4.3.6 of SSC rev 17. Read-write error recovery page. George Penokie stated that adding the test flag number to the report count field was a serious mistake. If he could have it changed he certainly would, but feels it is too late. In general George felt that the TapeAlert section was not well documented in that he felt it was difficult to understand from reading the SSC description. Dal Allan agreed that it was a mistake to not have developed this in the SCSI working group so a more conventional solution could have been developed. Dave Peterson noted that the documentation has been rewritten in revision 19, which was not been reviewed since it was not distributed until this meeting. He hoped that this would make the documentation more clear. Dal Allan and George Penokie felt that developing TapeAlert in isolation from the SCSI committee resulted in a non-optimal implementation, which did not correctly follow the SCSI architecture.

IBM Comment 98. TapeAlert requires that an error condition be reported on the next command in violation of the SPC, which says errors can be reported during any command. This issue was debated in detail last month in a discussion led by Rob Elliot, Compaq, in a different context. Rob had observed differences in the text between the SSC and the SPC. While Rob preferred the SSC text because it made testing easier, the decision last month was to remove the duplicated text from the SSC and conform to the SPC. This decision is also binding to the tape alert section. Stewart Wyatt was given an action item to feed this back to the TapeAlert group.

HP Comment 6. Proposed wording change about the physical model. Accepted

Seagate Comment 2. Requested subscription information for the SCSI (T10) reflector. Rejected

Seagate Comment 19. Definition comment. Accepted

6. FCP-2, T10 Working Drafts FCP2R02, Changes to FCP-2 T10/99-247r0: Bob Snively

Bob reported that he is making progress, but has not produced a new vision yet. Bob reviewed a marked up version of his change document T10-99/247r0, which he has not yet completed or distributed, but will be the rev 1.

TMF Logical Unit Reset made Mandatory. When questioned by Charles Binford, Rob Basham, argued that LUN Reset is better than a Target reset because it limits the scope of the reset appropriately. It is also required in FC-TAPE.

Bob believes that the Discovery work Dave Peterson has done is complete. It has not been circulated but will be included in the next revision of the FCP-2.

Name server extensions - This is Bob's assignment from another committee and he is working on it.

Clause 2.3, LSI Comment 008, Treatment of redundant name spaces. Rejected.

Default E_D_TOV value, 10 Seconds for fabric/point to point, defined in 99-145v0. Dave Baldwin took an action item to propose to FC-FS to change the point-to-point value of E_D_TOV to 2 seconds. Bob proposed that FCP-2 list the reference to where the default value is documented, identify how the default value could be changed and what the current default value is. There was some discussion about the appropriateness of including in the FCP-2 any references to time-outs defined in other documents. Bob felt that the issue needs to be summarized in FCP-2. The discussion disclosed the confusion about what the official values defined in FC-PH are.

Clause 11.2.7 will be retained in rev 3 of the FCP-2. The comment asking for it to be deleted was rejected.

Clause 3.2, LSI comment 002, text was added to clarify command and exchange termination.

Clause 3.4. LSI comment 004, Restriction of FCP_CONF usage. Bob is still opposed to this change (which was accepted by an overwhelming vote last month) and wanted a model justifying the change. Dal Allan noted he had provided an updated model on a floppy. Rob Basham clarified that the meaning of an initiator issuing the FCP_CONF is that the initiator won't ask for the response again - it will not be available as the resources will be cleared at the target. (The FCP_CONF cannot be sent blindly. The confirm comes from a higher level than an ACK does. The response must be validated before the confirmation is sent.) Charles Binford suggested that the text include a comment that the initiator cannot request the response again after sending a confirm. Bob made the comment that the confirm causes more command level overhead and increases bus congestion.

Clause 3.9, LSI Comment 10 accepted

Clause 3.11, LSI comment 012, Reject of retransmission requests. Bob feels that you need to try, but allows that there are cases where it cannot be done.

Clause 3.13, LSI comment 014, Bus inactivity clarification needs to be made consistent with SPC-2 (Also clause 3.14.)

LSI comment 016, Burst size during write clarification accepted

LSI comment 017, First burst size clarification accepted

Clause 3.17, LSI comment 018, RR_TOV default is too long. Bob thought that the default should be changed from a default of 300 seconds to a value that is derived from REC_TOV. A long discussion about RR_TOV followed. Bob thinks a much shorter minimum default would be desirable. Dale suggested that Bob make the change in the next revision for review.

Clause 3.18, LSI comment 019, Clarify FLOGI request accepted.

Clause 3.19, LSI comment 020, Support for Class 1 and 4 is not in FCP-2. Class 1 and 4 are acknowledged classes of service and should be treated the same as Class 2. The wording should refer to acknowledged and unacknowledged services rather than specific numbered classes of service. Bob wanted to prohibit mixed classes - someone noted that FC-PH already prohibits mixing acknowledged and unacknowledged classes of service. Bob said he would include the reference in FCP-2.

Clause 3.21, LSI comment 022, ACK(abort) to target. Noted that it was not covered in the FCP. Bob will look in the old FC-TAPE to see if it was defined. The target will probably have to send an ABTS.

Clause 3.24, LSI comment 025. Clause 11.2.3 in the FCP-2, FCP_XFER_RDY Recovery. The title of this clause only addresses writes. A new section is needed for read transfer ready. While Charles Binford saw this as an editorial change, Bob saw this as a serious technical omission. After some discussion a straw poll was held and the group voted unanimously to prohibit read transfer ready. The FCP-2 will state that the PRLI is required to disable read transfer ready.

Clause 3.28, LSI comment 029. Holding exchange information. Bob summarized the issues as (1) FCP_CONF should be allowed any time. (2) The FCP_RSP recovery may be rejected. (3) Chapter 11 should be reviewed for optional behaviors. A long discussion followed on queuing. Rob Basham wondered if knowing the Queue depth of the initiator would be helpful to the target. Charles Binford wants to use REC but not SRR and does not want to have to save the exchange information since he is not supporting SRR. Bob noted that the $n + 1$ requirement is wrong. Bob will review the problems and propose new text.

Clause 3.29 LSI comment 030. RO during recovery. RO needs to be defined more precisely for SRR recovery.

3.32 LSI comment 033, OX_ID reuse and RX_ID. If FCP_CONF is used RX_ID cannot be undefined.

HP Comments from Stewart Wyatt, requesting that recovery abort and ambiguous exchange definitions be included in the glossary. Bob rejected the request but agreed to edit the document to highlight the definitions.

Obsoleting recovery abort will be discussed next month.

Bob hopes to have a new revision of the FCP before for the next meeting.

7. Media Auxiliary Memory, T10/99-223r1 - Overview, T10/99-148r2 - New draft of proposal: Stewart Wyatt

The proposal has been moved to the SCSI working group meeting tomorrow for inclusion in the SPC-2. There will be no impact on the SSC document.

8. MCM Additions, T10/99-206r1: George Penokie

George presented an amended proposal. Bob wants to allow several lengths of this mode page, not to include the MCM stuff. He is particularly concerned about 0E length of the mode page. George argued that a shorter page could be returned without changing the document. Bob also wants it to be documented in an annex in case MCM is canceled. There were objectives to both of these requests. A one per company vote was taken. The motion was a recommend that the plenary vote to approve motion T10-206r1 as amended for acceptance into FCP-2. The results were 5 in favor, 3 against and 6 abstaining.

9. New Business:

- 9A. Partial Load/Unload T10/99-263r0: Paul Suhler

Paul presented a customer request for variants on the load and unload behavior. These changes were to support reading auxiliary memory in a tape cartridge. The changes were to allow loading the cartridge without threading the tape (called Load with Hold). And an unload variation instructing the drive to unthread the tape but not eject the cartridge (called Unload with Hold). Paul's suggestion was to add a HOLD bit to the reserved field of the Load/Unload commands. Some complications were brought up in the discussion. Eric Oetting asked about test unit ready - may need a new state to describe this condition. Model work is required including some new ASC/ASCQ. The appropriate place to document this capability was also discussed and it was suggested it may belong in the SPC-2. Paul will work on this proposal and post it to the reflector for additional discussion.

- 9B. EOM Clarification: Rob Basham

Rob has a customer request that EOM line up exactly for both reads and writes including subsequent mounts. He wanted confirmation that this request was unimplementable. The group unanimously agreed that this capability is never possible.

- 9C. PRLI - Queuing: Dale LaFollette

Dale relayed a request from an HBA vendor that wants targets to advertise whether they support queuing using PRLI. Bob Snively noted that the information is available in the inquiry command. Dal Allan noted that duplication of information will cause problems. (Duplicating inquiry data in the PRLI.) Group unanimously agreed that this was an unacceptable request.

9D. Multiple Fixed Block Read Recovery: Dale LaFollette

Dale led a discussion about how to deal with an error in an early sequence during a multiple sequence read operation. The subsequent sequences could be disrupted by the reception of a REC or a SRR and the transmission of an ACC. Dale noted that he might not be able to send the ACC until he completes transmitting the current sequence. Transmission of the ACC has to be on a sequence boundary. The consensus was that when the target transmits the ACC is its prerogative. The target is restricted from transmitting additional sequences following the error only after sending the ACC to the SRR and begins the error recovery operation. A long discussion about out-of-order followed. The issue was how an initiator can detect that a sequence that was transmitted after the error recovery started from one that was transmitted before. Charles suggests aborting sequences that were transmitted before. Eric noted that if the host cannot handle error recovery until the exchange is complete, the host can wait until after the response is received to start error recovery. Matt Wakeley noted that subsequent sequences may be aligned differently such that the data might line up differently in the sequences sent during error recovery than in the original sequences. Charles Binford noted that some HBAs can not have the data follow immediately follow the SRR ACC as some setup time is required. A solution, in this case, would be for the initiator to wait for the response from the target before initiating the error recovery. This discussion suggested that there may be two different approaches used by initiators to recover a read error depending on their turn time capability. The faster hosts will start recovery in the middle of the exchange with error. The slower hosts may wait for the exchange to complete before starting the error recovery.

9E. SMC-2 Project Proposal T10/99-218r0: Erich Oetting

Erich proposed a project to start an SMC-2 document. The goals would be to improve the operation with serial interconnects, handling interconnected changer devices and reporting error conditions. This working group agreed to forward the proposal to the plenary. The straw vote was approved unanimously.

9F. DSA, T10/99-226r0: Jim Coomes

Disable Soft Address discussion. Seagate's implementation always used just the hard address, not any of the other address types. The wording in the original documentation was vague on this. Jim noted that a new mode page bit could be assigned that supported the other views of this behavior or a vendor unique approach could be taken. The Tachyon chip implemented this so that it would also choose a previously acquired address also. This is the specific behavior that Seagate objected to. Dal Allan proposed softening the language, so that the requirement is that the hard address had to be used but could be selected in any of the previous initialization cycles. In addition the name should be changed to Require Hard Address (RHA). These suggestions were approved.

9G. FCP-2 Reset, T11/99-557v0: Charles Binford

This proposal was originally presented to FC-MI group and referred to this group for inclusion in the FCP-2. This proposal was to send a PRLO to those initiators with open exchanges that are terminated by another initiators actions. The purpose is to enable faster recovery. Bob Snively objected to the proposal, the PRLO blow away all information about the target - the device is reduced to an FC entity without target or initiator status. Ed Gardner proposed sending a FCP_RSP with a new response code (no SCSI status). Bob thought either a new ELS or ABTS would be more appropriate. The FCP_RSP approach has a problem with sequence initiative. ABTSL solves that problem.

Rob Basham became concerned during this discussion with the observation that a PRLI resets all mode pages and activity. If a new initiator does a PRLI with a target, which has completed a PRLI and started

activity with a previous initiator, will have all of the activities with the earlier initiator reset. He promised to take this discussion to reflector.

9. Next meeting Requirements Group

For the next T10 meeting in 2 months, the SSC should be finished, the FCP needs 4 hours, other activities will take between 2 and 4 hours.

10. Review new Action Items: Stewart Wyatt

Old items

#1 Carl Zeitler: Review need for process associators. - Due in October.

New items

#1 Stewart Wyatt, HP. The TapeAlert section must conform to the SPC error reporting requirements. (Any command not specifically the next command.) Report to TapeAlert group.

#2 Dave Baldwin, Emulex. Refer default E_D_TOV issue to FC_FS. Prefer a 2 second to 10 second for point to point connections.

#3 Paul Suhler. Proposal to add Hold addition to load/unload command. Updated proposal, post to reflector and continue discussion.

#4 Dale LaFollette, write up informative annex on read errors for the FCP-2.

#5 Eric Oetting, forward SMC-2 project proposal.

#6 Jim Coomes, DSA proposal updated: change name, soften requirement.

#7 Bob Snively, Revise RR_TOV default in next FCP-2 revision

#8 Bob Snively, FCP_XFER_RDY required to be disabled for reads in FCP-2.

#9 Bob Snively, Propose new text on requirements for targets to hold exchange information.

#10 Bob Snively, Include T10-206r1 in FCP-2

#11 Rob Basham, Investigate the effects on a PRLI from a new initiator on open exchanges from previous initiators over the reflector.

11. Adjournment: Group 7PM

Attendance List

Dale LaFollette	StorageTek	Stewart Wyatt	HP
Rob Basham	IBM	Paul Entzel	Quantum
Paul Suhler	Seagate	Edward A. Gardner	Ophidian Designs
Dave Baldwin	Emulex	Carl Zeitler	Compaq
John Artoux	JNI	Matt Wakeley	HP
Charles Binford	LSI	Bob Snively	Sun Microsystems
Jim Coomes	Seagate	Galen Fromm	Molex
George Penokie	IBM	Arlan Stone	UNISYS
Neil Wanamaker	Crossroads Systems	Jack Harwood	EMC
Rob Elliot	Compaq	Erich Oetting	StorageTek

David Peterson STK

Dal Allan

ENDL