

Joint T11/T10 FC-TAPE AdHoc Meeting  
February 10/11, 1999 Huntington Beach, California  
Stewart Wyatt, HP - Secretary

T11/99-104v0

1. Introductions. Dale LaFollette, StorageTek, the facilitator called the meeting to order just after 3 PM. As is customary, he had everyone introduce themselves.

2. Approval of this Agenda: T11/99-024v0 Group

The agenda was approved with the following changes:

Change Item 5. The latest draft of the profile is T11/99-069v0 (not 98-124vD)

Add to item 6, Recovery on Fabrics: Dave Ford.

Add to Item 7, Meeting requests: Group.

3. Approval of 12/16 Minutes: T11/98-619v0 Stewart Wyatt - Approved

4. Review of Old Action Items: Stewart Wyatt

Action items from Novembers T10 meetings

#1 Erich Oetting will propose to the T10 working group that the Write Exclusive Persistent Reservations be made illegal for SSC devices. Completed

#3 Erich Oetting will propose to the T10 working group that the communications devices command set be made obsolete. Completed

#4 Dave Peterson will attempt to document EOD behavior more precisely in the SSC. Dave has posted a new version of the SSC, deferred to the next T10 meeting

Action Items from Decembers T11 meetings

#1 Dave Peterson: Public/Private Loop authentication annex. Dave concluded this issue was adequately documented in the FLA. See discussion below.

5. FC-TAPE Document Review: T11/98-124vD (Corrected to T11/99-069v0) Dave Peterson

#1. SUN - Bob Snively, Comment # 20. "Remove clause 6.3 Loop Initialization features". This clause is labeled as "informational". The group agreed to remove it and insert a sentence that states that the initialization features required by FC-Tape will be defined by FC-AL2.

#2. Seagate - Paul Suhler, Comment #6 "Prohibit recover buffer data command". Currently the profile requires support of this command. Paul Suhler pointed out that it is not used in Seagate drives. Stewart Wyatt said that HP drives don't use it either. Dale asked if it could be retained as an allowed command. Bob Snively saw it as a feature that was used in older systems to recover data from the buffer when the drive ran out of media during a write. Modern systems have other means of avoiding this problem. Dale LaFollette thought it was useful for slow drives with large buffers with immediate command reporting. Bob argued that either every drive should implement the command or none should so that the host could use a common driver. Matt Wakeley, HP, argued that the host should keep a copy of the data in its buffer anyway so that this means of recovery should be unnecessary. The conclusion was to prohibit the recovered buffer data com-

mand, leaving an entry for it in the table, and add a comment that a host is required to keep a copy of the data until it is assured the data is safely on the media.

#3. LSI - Charles Binford, Comment #13, "Limit data overlay to error recovery". This comment was rejected in the last meeting in Charles absence. Since Charles was in attendance, the editor brought it up for review. Charles is concerned that with data overlay it is difficult for the host to tell that it has received all of the data. He was specifically worried that an entire sequence could be lost in a multiple sequence read operation. One solution would be to require continuously increasing sequence count. FC-PH requires this feature with streamed sequences, but the PLDA did not require it for the FCP\_RSP since existing initiators didn't support it.

During error recovery, the host sends an SRR requesting retransmission of the data. The current profile allows the target to begin the transfer at any point before the requested offset at the targets convenience. This makes it difficult for the host to ensure that all of the data has been received since current designs rely on counting the bytes received and comparing that to the amount requested. If the target sends more data than was requested, the existing host solutions become confused.

Dave Baldwin, Emulex, and Matt Wakeley, HP, agreed with Charles. Matt noted that he needs to allocate buffer space before sending the SRR and doesn't want the target to send more data than the space he has allocated.

Editor Dave Peterson saw these arguments as a violation of fundamental SCSI behavior. He felt that the hosts should use the relative offset of the buffer to check the data not the number of bytes transferred. He referred to the SCSI modify data pointer operation and was concerned that data integrity was being compromised by the existing host solutions.

Bob Kembel, Connectivity Solutions and others pointed out that parallel SCSI does not have a recover read data capability so that there is no precedence in SCSI for this situation.

Dal Allan, ENDL, noted that the issue boiled down to making life hard for targets or for initiators. He and others felt it was important to find a solution that could use existing hardware. The proposed solution was to require the target to begin error recovery on the precise boundary required by the host, assuming the target could begin recovery at some earlier boundary and bit bucket the extra data. This proposal was approved by a 8 to 1 vote.

The conclusion was to require targets to start recovery on the at the exact offset requested by the host. Targets are also required to use continuously increasing sequence count on streamed reads including the response. Initiators are expected to check for this functionality though the profile will be silent on this as some initiators are not currently capable of checking it. Finally data overlay will only be allowed in response to an SRR.

# 4 LSI - Charles Binford, Comment # 16 "wrong length for additional sense".

Charles became confused when he reviewed the PLDA requirements for FCP\_RSP when he saw what appeared to him as inconsistencies. These were explained by PLDA editor Bob Kimball who

was present.

Dale LaFollette thought the issue was over the buffer size required for the host to receive auto sense. Fibre Channel is different than parallel SCSI since it requires auto sense. The host has to allocate adequate buffer space for the auto sense data for every command it issues. The issue is the maximum FCP\_RSP frame that the host could expect to receive given the minimum frame buffer size. Dale notes that this was settled by this group long ago by asking targets how much sense data that they would send, less than 128. The minimum size has increased from 128 to 256.

Charles retracted his comment. Bob Kembel noted that for 256 byte frame the maximum sense space is 224 bytes. Dale and Charles agreed that restricting the size of the response frame is good and the current requirement is correct. The group rejected the comment, but table will be clarified anyway.

#5 Sun Bob Snively - Comment, "Add a data overlay bit for error recovery in PRLI."

Dave Baldwin, Emulex, wants a FC-Tape compliance bit in the PRLI.

Charles Binford wants CRN support bit in the inquiry data.

Dave Baldwin wants a FC-Tape compliance in the inquiry also.

Bob Snively, Dave Baldwin and George Penokie, IBM got into a long discussion about ports, targets, LUNs and whether this should be in mode pages, inquiry data or PRLI. George seemed to think something new needed to be invented. No resolution was reached.

#6 LSI - Charles Binford Comment # 25 "R\_A\_TOV instead of E\_D\_TOV". Charles noted that the profile incorrectly assumed that E\_D\_TOV would be used to time ELS replies. It should be 2 times R\_A\_TOV. Since R\_A\_TOV is specified by the Fabric, the value should be a reasonable.

Since the error recovery (clause 9 in previous drafts of the profile) is being moved to the FCP2 the issue was deferred to Bob Snively.

#7 Dave Ford, Clariion, had requested to make a presentation and was invited to do so. Dave presented several slides discussing the need for targets to recover resources from failed initiators in a fabric environment. His proposed solution was for targets to implement timeouts. His presentation was cut short when it was noted that the timeouts are already a requirement in the profile.

#8 Public loop authentication was discussed. This was the action item that Dave Peterson had from the last meeting to document the authentication process. On review of the FLA, Dave thought the issue was adequately documented. Questions from the participants could not find a need for additional documentation so the issue was dropped

#9 Connectivity Solutions - Bob Kimball Comment #15, "Requirement for loop port to recognize alias of x'00 00 AL\_PA". The comment was addressed in the FC-AL2 standard.

# 10 Connectivity Solutions - Bob Kimball Comment # 48, "FCP\_CMND Recovery"

Bob objected to the requirement for polling with REC at a specific rate. He felt that the host should not be forced to poll at exactly that rate. The group agreed to change the definition from a precise value to a minimum rate.

#11 Connectivity Solutions - Bob Kembel "Transfer of Sequence Initiative by SRR".

Bob Kembel sees this function as a fundamental architectural change. He thought a better solution would be a basic link service that changed initiative. Bob Snively agreed that it may cause problems for automated chips. Stewart Wyatt suggested that the change of sequence initiative could be indicated by a No-Op basic link service.

This issue was left open for additional thought and investigation.

#12 Connectivity Solutions - Bob Kembel Comment #56, "SCSI-specific behavior of SRR".

Bob thought that the SRR contained SCSI specific behavior and should have Type = x'08'. Since the FC-FS will be documenting all ELS he thought a more generic description might be needed. A discussion followed about whether SRR was indeed SCSI specific or could be used in other protocols. Bob Snively said he would look at the issue and would move the description to an annex in FCP2 for inclusion later in FC-FS.

At this point the group broke for the evening - Q-Logic was sponsoring a reception and food was available!

February 11, 1999, the meeting resumed at 8 AM. Dave continued the discussion of the letter ballots.

# 13 SUN- Bob Snively, Comment #27, "Incorrect value for ULP\_TOV".

Bob proposed that the value be changed to "Operation specific time plus two seconds". The current version of the specification is "Operation specific timer + 4 x REC\_TOV". This time needs to be large enough that recovery from an error can complete before this timer expires. The conclusion was to leave it as is. Clause 7.7 is to be changed by removing the second paragraph which is redundant and incorporating the note into the text.

#14 SUN - Bob Snively, Comment #26, "Unnecessarily strict requirement on NACA".

Bob thinks that implementing ACA will create a lot of trouble. He proposes changing note 3 in table 21 to read that ACA should be mandatory if tagged command queuing is used which is ordered or where state information queuing must be maintained. For other queuing environments it is recommended.

#15 SUN - Bob Snively, Comment #52, "Different second level error recovery"

Bob wants the same timeout to be used for ABTS, REC and SRR. ABTS is a BLS while REC and SRR are ELS. Dave noted that FC-PH is unclear about BLS timeouts. Use two times R\_A\_TOV

for all cases.

#16 SUN - Bob Snively, Comment #56, "Remove material provided in PLDA".

Dave had already removed the material. Some references will be added in FCP2.

#17 SUN - Bob Snively, Comment #57, "All sections are redundant with PLDA or other documents"

The editor has or will remove the material and will add text referencing the other documents.

#18 SUN - Bob Snively, Comment #58, "Clearing effects document s/b in FCP-2.

Accepted. The table will be moved out of FC-TAPE.

#19 SUN - Bob Snively, Comment #59, "APTPL should be required"

This would require tapes to save persistent reservations across power cycles. Dale noted that most tapes rewind and eject the tape when they are powered down. Bob asked if this was documented behavior. Bob is concerned that if the drive doesn't support persistent reservations that it needs to power up into a state where it has to be in an initialized state (i.e. not loaded) before it accepts any commands to support clustering requirements. Bob thought that this should be documented in the SSC (not in FC-TAPE). Deferred to SSC.

#20 SUN - Bob Snively, Comment #60, "Should TPRLO be made prohibited?"

Dave Baldwin noted that this is used today to clear third party reservations. Bob would prefer that persistent reservations be used. Deferred for Dave Baldwin's investigation.

#21 SUN - Bob Snively, Comment #61, " Reassignment of critical mode pages".

The SCSI committee has redefined the Fibre Channel mode pages. Bob also wants to consider which parameters should be invocable or prohibited. Bob proposed eliminating all of the options except EMDP and Maximum Burst Size. This was accepted.

#22 StorageTek, Comment #36. This proposal is to change the title of the profile to include media changers, "Fibre Channel Tape and Tape Media Changer (FC-TAPE). Accepted.

#23 LSI - Charles Binford, Comment #31, "SEQ\_CNT in ABTS accept frame".

Charles thought the Lowest SEQ\_CNT value does not consider all of the cases. Some members of the group thought that the profile should reference FC-PH and not include the material at all. Bob Kembel noted that the material is difficult to decipher in FC-PH and so was included in PLDA. This reference occurs in the error recovery clause that is being moved to FCP-2. Decided to refer to FC-PH

#24 Connectivity Solutions - Bob Kembel, Comment #29, "Response to selected link services".

Bob notes that Fibre Channel allows link services before PLOGI. PLDA and now FC-TAPE have prohibited any link services before completed login. Bob suggested some cases where allowing some link services before login is completed. One example is error recovery on a failed login. Deferred.

6. T11 New Business: Group

7. T10 New Business: Group

7.1 SSC-2 Dave Peterson

7.2 FCP-2 Bob Snively

New business

Goal was to complete review of comments by April. The editor suggested he would like to get as much time as he could get. Bob Snively wants some time for FCP-2. There is also some SSC business to complete as well

8. Review New Action Items: Stewart Wyatt

#1. Dave Peterson/Bob Snively - Update the error recovery documentation to resolve the timeout for ELS.

#2. Group - Resolve the issue of signalling support for FC-Tape error recovery, CRN support and data overlay in either/or inquiry data, mode pages, PRLI or by some other means.

#3. Group - Consider the effects of transferring sequence initiative by means of the SRR. Is there a better way?

#4. Bob Snively - Consider whether SRR is SCSI specific. But in an annex in FCP2 for inclusion later in the FC-FS.

#5. Dave Peterson - Report the omission of BLS timeout in FC-PH.

#6. Dave Peterson - Consider modifying the tape model to require that the tape power up uninitialized, unloaded state. SSC

#7. Dave Baldwin - Investigate eliminating TPRLO.

#8 Dave Peterson/group - Resolution of need for authentication documentation.

9. Adjournment: Group