

Joint T10/T11.3 Activity Working Group AdHoc Meeting T11/99-754v0  
December 7, 1999, Reno Nevada

Stewart Wyatt – Secretary

1 Introductions: Dale LaFollette

The meeting started at 1PM. Dale had the group introduce themselves.

2 Approve this Agenda: T11/99-725v0 Dale LaFollette

Approved with minor additions.

3 Approve 11/02/99 Minutes: T11/99-679v0 Stewart Wyatt

Approved

4 Review Old Action Items: Stewart Wyatt

#1. Charles Binford, LSI Logic. Refer default E\_D\_TOV issue to FC\_FS. Prefer a 2 second to 10 second for point to point connections. Reassigned to Dave Peterson

#2. Charles Binford, LSI Logic. Propose a new RESP\_CODE 0x06: Command cleared by another initiator. Take new status code to SAM-2. Ongoing

New Action Items

#1 Dale LaFollette, STK. Prepare agenda for the January T10 meeting in Australia - Ongoing

#2 Bob Snively, SUN Microsystems. Facilitate and take minutes for the January 2000 Joint Activity meeting in Australia. - Ongoing

#3 Bob Snively, SUN Microsystems. Investigate whether it is appropriate to include the discovery ELS (RTIN and RNID) in the list of ELS to be accepted before completing login. – Moved to FC-FS

#4 Bob Snively, SUN Microsystems. Move the diagrams of Annex C into clause 11, making them a normative example. – Next FCP revision

#5 Bob Snively, SUN Microsystems. Verify the multiple of 4 byte block length change is acceptable for both the SSC-2 and FCP-2 by posting a proposal to the reflector and bringing it up at the SCSI plenary. – Plenary presentation completed but not the reflector item

#6 Dave Peterson and Dal Allan. Review the current FCP-2 for limitations to out-of-order error recovery implementations. – Carl Zeitler presentation this meeting, additional action items continue the issue

#7. Bob Snively. Identify the problems with allowing unlimited process associators and prepare a review. – Bob made a counter proposal presented over the reflector and at this meeting

#8 Dave Peterson. Review Annex E SCSI Device Discovery Procedure, to determine what the differences are for target and initiator requirements are. - Completed

+++ Joint T10/T11.3 +++

5 FCP-2 T10 working drafts FCP2R03 and Change comments T10/99-325r1 Bob Snively

Bob reviewed the new revision of his change document. There were only a few issues that generated extended discussion.

Out of Order issue: Bob invited Carl Zeitler to review his presentation. He noted that the in-order requirement compromises complex fabrics. Stewart Wyatt observed that in a backup environment, allowing out-of-order would require additional buffering for tape drives which operate sequentially. He noted that out-of-order requirements would reduce the costs of fabrics while increasing the cost of peripherals.

Carl noted that time out values are critical in out-of-order error recovery. Currently REC\_TOV is defined as ED\_TOV + 1 second. With these values in an out-of-order environment, the target could receive a REC before it received the command. The solution is to make the REC\_TOV large enough to guarantee that the frames are not still in the fabric.

Carl presented a series of overheads that used ladder diagrams to demonstrate the various error recovery cases for both Class 2 and 3. The first set for commands, transfer readies and responses were relatively straight forward. There was some confusion over Carl's use of ABTS. Carl was using the ABTS to recover status on the sequence and set up a recovery qualifier, not to abort the exchange. In Class 3 Carl used RA\_TOV/2 for the timeout value instead of REC. Carl questioned the value of REC in Class 2.

The data transfer examples had more complexity. Some issues were raised. Carl's overheads indicated the detection of a sequence error when the frame was lost which is only true with in-order error detection. Out-of-order does not detect an error until the timeout expires. The overheads also showed ABTS being used incorrectly. Another concern was expressed over the ability of an HBA to detect that a frame has been transmitted to start a timer since these events occur at different levels in the HBA. Matt Wakeley noted that ED\_TOV will expire at the sequence recipient before the sequence initiator will detect that an error occurred.

There was a lot of discussion about recovering lost ACKs in Class 2. In an ACK 1 form, the only ACK that must be recovered is the last one of the sequence. An RRQ is always required to avoid confusion if the missing ACK pops out later.

Carl also raised the issue of aborting a failing sequence in Class 3 during an exchange that is transferring a large amount of data. This group had previously discussed this issue and the feeling was that with the high transfer rate, the delay would not be severe enough to justify action. Nevertheless, Dale LaFollette has a proposal to use SRR to get the target's attention that will be an annex in the next revision of the FCP.

After considerable discussion, Dale LaFollette asked the group what they thought. Dal Allan and Ed Gardiner both thought the issue was very manageable, being largely an issue of timeouts, and thought that it should be included. Bob Snively thought it would require more extensive code structures in the host and be less robust. Dale asked Bob how long he thought it would delay the FCP-2 to add it. After some discussion, the answer was six months.

Stewart Wyatt expressed concern that delaying the FCP would reduce support for developing drivers and applications for use by early implementations, which will only support in-order delivery. He was concerned that this would delay the market for Fibre Channel tapes. Stewart was also concerned about making a clear distinction in the documentation between in-order and out-of-order systems. Jim Coomes asked how an in-order and out-of-order peripheral would be distinguished. A means will need to be identified.

Stewart proposed completing the FCP-2 as is and starting a FCP-3 that supports out-of-order. George Penokie seconded the motion. After further discussion, the motion was withdrawn to allow Bob and others to better access the work involved and define the impact.

Process Associators: Bob Snively had created a counter proposal to replace the need for process associators. Bob noted that targets could use LUNs for separate representations. Other groups are developing access control. The only missing piece is providing a means of placing multiple initiators behind a single port. The solution Bob is proposing is to have a separate ID with a distinct world wide name for each initiator. He thought that his approach would have the least impact on switches.

Jeff Stai expressed a concern that the scope of this discussion be strictly limited to FCP. (FS has a separate agenda for process associators.)

While the proposal was positively received, Dal Allan and Matt Wakeley questioned the number of addresses available to define the initiators. Dal noted that ten years ago 8 bytes was considered necessary. Carl Zeitler thought that a single byte (256 initiators) would be adequate. Later Carl expressed concern about the numbers required for Infiniband support. Ed Gardner reassured Carl that these numbers were adequate.

Dal thought each application could require a unique initiator to support zoning, assuming that each application would need its own zone. Bob Snively argued that process associators had not been involved with zoning.

Finally the group came to a conclusion. Bob's proposal would be the solution for FCP. (Process Associators would be a FC-FS function to appear in FC-SW-2.) The policy will be that Process Associators does not participate in task identification. It was suggested that the group endorse this decision. The lone dissenter was Matt Wakeley over concerns that the available number was too small. (This concern may provide an incentive elsewhere to develop header to increase the address space – and these may look like process associators!)

Incorrect use of recovery abort: Bob noted that Charles Binford's proposal (discussed below) reduces the number of ambiguous changes and the incorrect usage of recovery aborts. A mode bit enables the features of Charles' proposal. Bob thought that this function should be required. Bob expects that ABTS will be used to cancel any remaining ambiguous exchanges, even requiring this functionality from a Class 3 target.

EMDP bit: Bob had required the EMDP bit to be set for SRR error recovery. Matt Wakeley expressed concern that setting this bit implies support of out-of-order delivery, which is not what was intended. Charles Binford proposed that REC/SRR error recovery be allowed without setting this bit. This proposal was accepted.

#### 6 FCP-2/FC-GS-3 Update T11/99-710v0 Dave Peterson

Dave has made a proposal that was accepted by FC-GS-3 working group to assign four bits for FC-4 specific usage. The bits will indicate whether the device supports SCSI target and/or initiator behavior. There was some discussion about this capability.

#### 7 REC/SRR Questions T11/99-728v0 Neil Wanamaker

Neil's slides were reviewed during FCP-2 discussion. 1<sup>st</sup> and 3<sup>rd</sup> slides accepted. The 2<sup>nd</sup> slide needs review to resolve an ordering issue.

#### 8 Task Management Questions Dale LaFollette

Dale noted that the FCP-2 ambiguously describes the targets response to a task management function as either accepted or performed. Ed Gardiner, referring to SAM, noted that these terms are the same. Dale asked how a device should respond before completing a target reset. The answer is to respond with a busy.

+++ T10 +++

#### 9 SSC: T10 Working Drafts SSC-R22 and change document T10/99-228r4 Dave Peterson

No news, waiting for the public review to close.

+++ New Business +++

#### 10 CMDS Cleared 99-722v0 Charles Binford

This proposal was originally made to this group. Charles was sent to present it to the SCSI working group. This agenda item was an update. Charles made a presentation to the SCSI group, which was positively received. He is planning on working with Ralph Weber to make a formal proposal.

Charles has two objectives in this proposal: The first is to notify other initiators that their exchanges were cleared by a TMF from another initiator. The second is to quickly clean up aborted I/Os to remove ambiguous exchanges.

The solution is to introduce a new SCSI status 'CMD Cleared', sent to all initiators with outstanding aborted commands including the one who cleared them. No ordering requirements are assumed. Other causes of clearing commands (PLOGI, LOGO, etc.) are unchanged. Charles prefers that targets send a PRLO to all initiators if it receives a TPRLO. The Control mode page will enable this ability.

#### 11 Next Meeting Requirements: Dale LaFollette

Dale asked how much time would be required for the meeting in two months. Dave Peterson was asked if he needed time for reviewing the SSC after the public review is completed. He didn't think so. Bob Snively hopes to have completed the letter ballot on the next revision of FCP-2 by then and requested three hours minimum.

#### 12 Review New Action Items: Stewart Wyatt

#1. Dave Peterson. "Reasonable" timer values proposal, including Dave Baldwin's proposed reduction of E\_D\_TOV.

#2. Charles Binford. Command cleared proposal for the SCSI working group.

#3. Dale LaFollette. Prepare agenda for January meeting in Australia.

#4. Bob Snively. Facilitate and provide minutes for the January meeting in Australia.

#5. Bob Snively. Move the diagrams of Annex C into clause 11 making them a normative example.

#6. Bob Snively. Post the four byte multiple fixed block length and error recovery decision to the reflector.

#7 Bob Snively. Include revised Annex E (T11/99-340v3) in the next revision of the FCP-2.

#8 Bob Snively, Carl Zeitler, Dave Peterson: Review the impact of adding out-of-order delivery to the FCP-2. Schedule a review in the February meeting. Have Bill Martin review the proposal.

13 Adjournment: The group adjourned at 5:40 PM

Attendance List:

Dale LaFollette	StorageTek	Stewart Wyatt	HP
David Peterson	STK	Roger Cummings	DPT
Joe Breher	Exabyte	Neil Wanamaker	Crossroads Systems
John Scheible	IBM	George Penokie	IBM
Bob Snively	SUN	Vit Novak	SUN
Damian Bannon	SSL	Pak Seto	Quantum
Chuck McKnett	JNI	Steve O'Neil	CMD
Matt Wakeley	HP/Agilent	Bill Martin	Gadzoox Networks
Jim Coomes	Seagate	Charles Binford	LSI
Danny Ybarra	TI	Carl Zeitler	Compaq
Gene Milligan	Seagate	Arlan Stone	Unisys
Dave Guss	TI	Mark Hamel	Compaq
Paul Suhler	Seagate		