

Joint T11.3/T10 Activity Working AdHoc Group
Huntington Beach, CA
February 16, 2000

T11/00-118v1
Revised March 6, 2000

Stewart Wyatt, HP, Secretary

1. Introductions: Dale LaFollette

Facilitator Dale LaFollette called the meeting to order and had the participants introduce themselves.

2. Approve this agenda: T11/00-003v1 – Dale LaFollette

The agenda was approved but the items were reordered to postpone resolving the FCP-2 letter ballot comments until all of the other agenda items had been completed.

3. Approve 01/13/00 minutes: T10/00: T10/00-130r0 – Bob Snively

Stewart Wyatt reviewed the minutes of last months meeting, which were approved without comment. Stewart also thanked Bob Snively for his taking those minutes.

4. Review old action items:

#1. Dave Peterson: Propose reasonable timer values, including Dave Baldwin's proposed reduction of E_D_TOV. Completed.

#2. Charles Binford: Proposal to notify initiators of cleared commands to be taken to the SCSI working group. In progress.

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

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

#5. Bob Snively: Move the diagrams of Annex C into clause 11, making them a normative example. Elected not to make the change for editorial reasons.

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

#7. Bob Snively: Include revised Annex E (T11/99-340v3) in the next revision of the FCP-2. Completed in FCP-2 revision 4. This item is still open and the subject of a letter ballot comment.

#8. Bob Snively, Carl Zeitler, and Dave Peterson: Review the impact of adding out-of-order delivery to the FCP-2. Schedule a review in the February meeting. Carl Zeitler made the presentation, which was accepted.

#9. Paul Suhler will prepare a proposal for addressing the problems in the LOCATE and READ POSITION commands for consideration at the February meeting. Completed.

#10. Bob Snively will include a comment requesting the removal of Annex J in the Sun letter ballot comments. Completed

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

5. New Business

5 A. Carl Zeitler FCP-2 Out Of Order Presentation T10/00-137r0

Carl duplicated the error recovery diagram examples in Annex D of FCP-2 Rev 4 and updated each of the cases to accommodate out-of-order frame reception. Carl stated an assumption of his that performance is not an issue with error recovery.

Carl started with Class 3 examples, commenting that the changes to Class 3 were relatively minor. Carl provided a new definition of REC_TOV, which was equal to half of R_A_TOV. To briefly summarize the Class 3 examples, there are two significant cases. The first is where a frame was assumed lost on a “one way” trip. Error recovery is initiated after REC_TOV. An example would be a lost data frame on a write. The initiator waits REC_TOV then sends an REC. The second case is when the frame was assumed to be lost on a “round trip” then two REC_TOV waits (equal to R_A_TOV) are required. An example occurs when a FCP_CMND only requires a FCP_RSP and the FCP_RSP is lost. The initiator waits REC_TOV before sending the REC, and then waits another REC_TOV for the original response before initiating error recovery.

Matt Wakeley noted that there is no requirement for when the next sequence is sent in FC-PH. The recovery timers assume that the reply is sent instantaneously and the delays are all in the fabric, but that is not a requirement of the standard.

Carl noted that Class 2 requires more changes than Class 3. He also claimed he was using Class 2 in the “classic” sense. Carl’s approach followed the procedure described in FC-PH. It was soon apparent that the participants were not conversant with the standard procedure and a lot of discussion and references to FC-PH occurred. Carl admitted that he had to do some reviewing himself. Carl defended the common error recovery approach between the classes of service already established by this committee. His proposal requires ABTS to abort sequences (not exchanges) and requires the use of recovery qualifiers.

One issue that surfaced was that Sequence ID is not a qualifier of Sequence Count. In recovering from an error, a new Sequence ID must be used and the Sequence Count must be incremented from the value used in the ABTS. Bob Snively noted that this should have been required in the in-order case as well.

Carl observed that the RRQ was optional for the target when an ACK was lost. This prompted a long discussion. Bob Snively thought that the RRQ was necessary to prevent the Sequence Recipient from reusing the OX-ID of the lost ACK. It was resolved that the Sequence Initiator must retire the RX-ID until the RRQ expires, but that there is no obligation on the Sequence Recipient. Matt Wakeley was concerned that without establishing an RRQ that the EE_Credit count would be corrupted if the ACK did appear later. A resolution of this dilemma was to recover the credit when the RRQ was established and not claim credit it for the ACK if it does appear later. There was considerable discussion over these issues and numerous references back to FC-PH.

Another issue that excited considerable discussion was whether a BA_RJT is allowed for an ABTS. Carl noted a special case of a lost command. If the first frame of an exchange is lost, the Exchange Responder cannot reject the ABTS even though it has unfamiliar with the Exchange. The ABTS must be accepted to set up the RRQ in case the command appears later. There was some concern that issuing a BA_RJT needs to be prohibited since it does not set up a recovery qualifier.

After Carl had presented his slides the group reviewed the impact. Bob Kembel noted that the FCP-2 error recovery proposal allows terminating Exchanges in a manner not documented in FC-PH. Stewart Wyatt asked how a device that only supports in-order command reception would signal that capability. The answer is that in-order requirements are given to the fabric in the FLOGI. The Class 3 case requires the changes to the SEQ_CNT and SEQ_ID, noted earlier, though that change is independent of the in-order issue. Class 3 also requires changing REC_TOV.

A motion was forwarded for the groups approval by Carl Zeitler and seconded by Bob Snively: Do we want to use Carl's proposal as amended (which supports both in and out of order delivery) as the only error recovery procedure for FCP-2? On a company vote the measure passed 14 for, 3 against, 2 abstaining.

5 B. Data Transfer Integrity T11/00-021v0 Dale LaFollette

Dale LaFollette made a presentation that he had prepared with Matt Wakeley. Dale observed that the PCI bus used internally to some initiators and targets is the weak link in protecting the data during a transfer since there is only 1 bit of parity covering 32 bits of data. (Ed Gardiner noted that some designs do not even check the parity.) Dale's proposal included several approaches to protecting data over the entire transfer but not including the actual storage media.

The most promising proposal was to have the target create a checksum, which it returns to the host in the response. The host also creates a checksum, which it compares to ensure that the data transferred without error. This approach works well for serial protocols like Fibre Channel, but not for parallel protocols.

George Penokie vigorously opposed the proposal, noting that "the paranoid have already solved this problem". His reference was to vendor unique solutions and he felt that there is no need for a standards solution. Bob Snively countered that more people are becoming paranoid. Joe Breher expressed concern with the protocol attempting to solve internal device problems. Ed Gardiner expressed concern about obtaining host side support for this feature.

Dale called for a straw poll, "should this proposal be continued?" The results were 6 for, 5 against with a large number abstaining.

5 C. Read FCP_XFER_RDY T11/00-067v0 Dave Peterson

Dave's proposal was to add a FCP_XFER_RDY from the initiator for reads. FCP had a transfer ready for reads that came from the target, which was obsoleted in FCP-2. Normally the host sends commands for space that it has available and does not need this level of flow control. Dave justified his proposal by noting that some times an intermediate device such as a router does not know the available buffer space in the host and wants to use the transfer ready for flow control. This proposal led to a long discussion. The longest argument about target's role in controlling data transfers. Charles Binford noted some precedents for this type of behavior. Bob Snively expressed concern about how it would affect error recovery. There was also some paranoia expressed about the effect this would have on performance and the risk of deadlock. Finally Dale LaFollette called a straw poll asking if the group would support continued effort to develop this proposal. The results were for continuing 5, against continuing 7, abstaining 6.

Matt Wakeley expressed frustration with those who voted against further investigation of optional proposals, referring to this proposal and the previous one. The reluctance comes partly from target manufacturers who fear that optional features would become mandatory for them.

+++ T11.3 +++

6. T11.3 New Business:

None.

+++ T10 +++

7. T10 New Business

7 A. SSC Public Review Comment

An SSC public review comment was received from Paul Suhler requesting restoring the SET CAPACITY Command, which was in early versions of SSC but not in the final version. There was some procedural discussion between George Penokie and Dave Peterson. As the discussion continued the history of this command became more suspect. Checks in other documents indicated that the command had never actually existed. The conclusion was that it was only a proposal that had never been finished. Paul wanted to have the command revived to support reducing the capacity of a large tape. He promised to consider using a partition size or some mode page as an alternative to using this command

7 B. Response to PRLI/ACC with both I/T bits set. Neil Wanamaker

This issue had surfaced in the FC_MI group. During interoperability testing it was found that some hosts and targets will fail to complete PRLI if the other party sets both the host and target bits. Examples of this case include disks implementing the XOR command and tapes implementing the COPY command. A specific issue is when a target, which can

also function as an initiator, attempts to send a PRLI to an initiator. Some initiators will not complete the PRLI unless they initiated it. Neal is working on the appropriate response in each of these cases.

7 C. SSC Large Block Addresses T10/00-135r0 Paul Suhler

Paul noted that with the large capacity of emerging tape drives there is a need to increase address space in commands. Of the methods proposed, using 16 byte commands appears most acceptable. Paul marked up his proposal with inputs from the group, which included the affected commands.

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

8. FCP-2: T10 Working Drafts FCP2R04 and Letter Ballot Comments T10/00-150r0 Bob Snively

The letter ballot review had been put off until last. Bob noted that there were 1350 letter ballot comments to review. Bob said he planned to go through them and answer the editorial comments and the obvious technical comments. He would review the more difficult or controversial comments with the group. The attendance had diminished as the evening wore on. A reception was in progress in a near by room. As it was quite late, the meeting was adjourned without starting the letter ballot review.

9. Next Meeting Requirements – Dale LaFollette

The next meeting is being held in Dallas, Texas, at the Crowne Plaza Suites sponsored by Texas Instruments. The main agenda item will be reviewing the FCP-2 letter ballot comments. The meeting was originally scheduled for Tuesday March 7, which conflicted with the Parallel SCSI Working Group. After some discussion the meeting time was changed to Monday March 6. The meeting will start at 9 AM and is scheduled to run until 6 PM.

10. Review New Action Items: Stewart Wyatt

#1 Carl Zeitler – Update the Out-Of-Order proposal drawings in T10/00-137r0.

#2 Carl Zeitler – Determine the correct response to a lost ACK: BA_RJT or BA_ACC.

#3 Bob Snively – End exchange cases in FC-PH does not include the class 3 case of lost FCP_CONF. Needs to be added to FC-FS. Check for other cases.

#5 Paul Suhler – SSC letter ballot comment requesting restoring the Set Capacity command. Paul will investigate some other means of limiting the available capacity of a large tape.

#6 Neil Wanamker – Revise proposal defining behavior with both target and initiator bits set in PRLI.

#7 Paul Suhler – Revise proposal for Large Block Addresses in the SSC.

Old Action Items

#8. Charles Binford: Proposal to notify initiators of cleared commands to be taken to the SCSI working group.

#9. Bob Snively: Include revised Annex E (T11/99-340v3) in the next revision of the FCP-2. Completed in FCP-2 revision 4. This item is still open and the subject of a letter ballot comment.

#10. Bill Martin: Bob Snively requests that Bill Martin review the Out-Of-Order proposal looking for corner case problems.

11. Adjournment: Dale LaFollette

Attendance:

Dale LaFollette	STK	Stewart Wyatt	HP
David Peterson	STK	Arlan Stone	Unisys
George Penokie	IBM	Ralph Weber	ENDL
John Lohmeyer	LSI Logic	Matt Wakeley	HP/Agilent Tech
Dennis Moore	KnowledgeTek	Charles Monia	Adaptec
Neil Edmunds	Xyratex	Predraig Spasic	HP
Paul Suhler	Seagate	Horst Truestedt	TrueFocus
Jim Coomes	Seagate	Curt Ridgeway	LSI Logic
Bob Snively	SUN	Edward A. Gardner	Ophidian Designs
Robert Reynolds	Crossroads Systems	Pak Seto	Quantum
Bret Ketchum	CNT	Kumar Malavalli	Brocade
Bob Kembel	Connectivity Sol.	Craig Stuber	JNI
Mike Fitzpatrick	Fujitsu	Jeff Stai	Q Logic
Stephen O'Neil	CMD	Joe Breher	Exabyte