

Joint T10/T11.3 Activity Working Group AdHoc Meeting T11/99-620v0
October 5, 1999 - Fort Lauderdale, Florida

1. Introductions: Facilitator Dale LaFollette called the meeting to order at 10:30 AM and had the participants introduce themselves. The previous meeting ended early, enabling this meeting to start earlier than the scheduled 1 PM start time
2. Approval of this agenda: T11/99-595v0 - Group. Approved
3. Approval of 9/14/99 Minutes: T11/99-564v0 - Stewart Wyatt. Approved
4. Review old action items: Stewart Wyatt

Old items for the last meeting

#1 Carl Zeitler: Review need for process associators. - Due in October. Carl felt that process associators should be retained and cited support from IBM though he did not have a formal justification. Dal Allan supported Carl's position, noting that he felt process associators could be useful in a standard means of providing zoning in a fabric. This topic was discussed later in agenda item 10C.

New items for the last meeting

#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. Stewart sent a detailed email to Stephen Gold, however Stephen was traveling and had not replied before this meeting.

#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. Deferred, as Dave was not in attendance.

#3 Paul Suhler. Proposal to add Hold addition to load/unload command. Updated proposal, post to reflector and continue discussion. Deferred, as Paul was not in attendance.

#4 Dale LaFollette, write up informative annex on read errors for the FCP-2. Dale was waiting for a template from Bob Snively before completing this item.

#5 Eric Oetting, forward SMC-2 project proposal. Completed per John Lohmeyer.

#6 Jim Coomes, DSA proposal updated: change name, soften requirements. Jim was not in attendance but his proposal had been made and reviewed over the reflector. The reflector traffic was reviewed. There was some discussion about rewording the proposal to make it not binding on fabric attach topologies where it is not meaningful.

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

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

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

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

#11 Rob Basham, Investigate the effects of a PRLI from a new initiator on open exchanges from previous initiators over the reflector. This was discussed and resolved in the meeting (agenda item 10A), though Rob should review the results.

5. SSC: T10 Working Drafts SSC-R20 and T10/99-228r2 - Dave Peterson

Dave brought up several of Gene Milligan's comments for discussion.

GEM#30 Gene noted that in Clause 5.2.1, the text requires that "the logical unit is not ready ... whenever all medium access commands report CHECK CONDITION status". "All" should be changed to "any".

GEM#38 Referring to Clause 5.2.8, Gene asked what does "unique for the medium" mean. Dave said he would look into a definition.

GEM #39 Gene asked for more a clearer statement defining sequential device positioning. Dave said he would try and clarify the text.

GEM#50 Clause 5.2.12.2 Gene noticed a difference in behavior in clearing flags internally. The difference is intentional - no change.

Dave reported that he was nearly completed with the SSC.

Dale LaFollette asked who would be attending the January 2000 T10 meeting in Australia. Ten people responded out of the 33 who signed the attendance sheet.

6. FCP-2: FCP2R03 and T10/99-247r1 - T10/99-247r1 - Bob Snively

George Penokie had a question about loop initialization. He noted that during loop initialization, the device is required to check its hard address. A dual port disk drive has a shared ESI port that is also used to access the hard address. If the ESI port is busy when the loop initialization request comes, the port is blocked from reading the hard address. Bob's reply was the drive should free up the ESI port as soon as possible to meet the requirement. In general the group thought that there was adequate time to complete the ESI service and read the address within the allocated time.

Bob had completed a new revision of the FCP-2 late last week (FCP2R03). He had also updated his change document, "FCP-2 changes from revision 02 to 03" (T10/99-247r1) Bob reviewed the change document. These minutes note the clauses in the change document where there was significant discussion.

After reviewing Clause 3.4, Bob noted that FCP_CONF was being restricted to actual SCSI status responses. It is not a valid request for Task Management Functions or for Fibre Channel responses that do not include valid status. Charles Binford asked how to identify a response without valid SCSI status. The "first cut" proposal was to make the RSP_CODE and the SCSI Status mutually exclusive. If this is not possible a new technique will be devised.

Clause 3.11, Additional issue 2: SRR ACC timing. Bob specified that the SRR ACC should meet the FC-PH ELS timing requirements. Charles Binford asked if he could be assured that the ACC would arrive before the recovery data. Bob said no, the data could arrive before the ACC. After some discussion the following requirement was added that the SRR accept must be transmitted before the target initiates any error recovery and that the target cannot send any pre-error recovery data after transmitting the accept.

Bob had forgotten a previous agreement that the targets would only transmit the data that had been requested by the initiator. This may require the target to begin recovery at an earlier point if that is necessary to access the requested data and discard the data up to the requested start point. Several corrections to the text are required to correct this oversight.

Bob noted two sections in the review document, clause 3.25, Permission to do recovery, and clause 3.28, Hold exchange information, and asked the group to review these areas and clause 11 in the latest FCP-2 draft on error recovery.

Clause 3.29 was one of the places where Bob erroneously did not require the target to start error recovery on the required boundaries. The last paragraph needs to be removed. Similar errors are in clause 3.30 and

3.31. Dale LaFollette noted that the error recovery boundary must be modulo 4 relative offset. Bob promised to add that to the next revision

In referring to clause 3.33, Bob asked why the S_ID appears in the payload of the REC, since it is duplicated from the header? Roger Cummings thought there might be some historical reason. Bob Snively said he was going to propose the field be reserved.

Referring to HP comment in clause 4.14, Bob thought the arrows in the drawing were correct. Stewart Wyatt checked the document. The conclusion was that the arrows are correct, but Bob will check the proximity of the labels.

Bob told the group that next month he wants to reconsider the in-order requirement.

The process associator issue is still open and may be a pacing item in completing the FCP-2. Bob proposed restricting the processor associators image to one per pair of SCSI devices which he thought might be an acceptable approach. He promised to send out notification and sample text for review.

Bob asked the group to review clause 6.3, Incorrect use of Recovery Abort and FCP-2 clause 11.5.1.

In clause 6.4, Specification of formats for ELS, Bob proposes moving FCP-2 clause 11.7 to an informative annex.

In clause 6.5, Multi-Initiator, Bob proposes moving FCP-2 clause 11.9 to an informative annex.

Bob hopes to have FCP-2 Revision 4 ready two weeks previous to the November for a letter ballot vote. It is unlikely that the process associator issue will be resolved in time. If not, the letter ballot vote will occur in January.

7. Tape Connector 8072_011.txt - Stewart Wyatt

Stewart Wyatt reported on the progress made on the Fibre Channel tape connector at last month's Small Form Factor committee meeting. When the proposal was made to use the 80 pin SCA-2 connector, it was thought that mixing the SCSI and Fibre Channel devices and back planes would not result in damage. Later it was discovered that the SCSI signal positions that were used for additional 5 Volt power in this Fibre Channel for had a maximum spec of 4.1 Volts for LVD or active negation single ended pads. The decision in the SFF was to assume that the form factor would preclude mixing the applications. All of the SCSI SCA-2 80 pin devices are 3 1/2 inch form factor devices and this Fibre Channel tape proposal has been modified to restrict it to 5 1/4 inch form factor devices. When this change was approved, additional 12 V lines were added to the proposal at Seagate Tape groups request. These positions are also used for SCSI signals in the SCSI standard.

George Penokie expressed concern about restricting SCSI disk drives to the 5 1/4 inch form factor. Gene Milligan said that the Seagate 5 1/4 inch drives never used the SCA-2 connector. No one was aware of any other 5 1/4 inch SCSI disk drives.

8. DSA/RHA: T10/99-226r2 - Jim Coomes

During a brief discussion of this proposal, there was a question about the affect this bit would have on direct fabric attach topologies. The conclusion was that this bit was irrelevant to a direct fabric attach device and would have no effect. The wording of the proposal may need to be clarified.

9. Command Cleared Notification: T10/99-591v0 - Charles Binford

Charles goal is that when an initiator clears tasks in a target for other initiators, the target has a means of informing the other initiators it has open exchanges with to reduce the recovery time. Last month Charles proposed using PRLO as the means of notification. When this approach was reviewed it was found to be

unacceptable. The alternative is a notification for each outstanding I/O. One approach that was proposed was to use ABTS since it avoids the issue of sequence initiative. Charles felt that using ABTS would result in an ugly implementation.

Charles new proposal is to define a new RESP_CODE 0x06: Command cleared by another initiator. The target would have to wait for sequence initiative to transmit the response, but that is not expected to be a problem.

Bob Snively thought that this would have applications outside of Fibre Channel and could be used in other SCSI applications. Dave Peterson preferred the ABTS approach since he thought that ABTS will be useful in Class 2 and out-of-order delivery systems.

Charles proposal was taken to a straw poll vote. The results were 7 in favor, 0 opposed.

Action item: Take new status code to SAM-2

10. New Business

10A. FCP-2 Clearing Effects - George Penokie

This discussion followed from last months observation by Rob Basham that a PRLI would cause exchanges in progress from previously logged in initiators to be aborted. The PRLI would also reset Mode pages to their power on value. Bob Snively noted that a PRLO changes the status of a node from a SCSI target or initiator to an undefined Fibre Channel node unless an implicit PRLI is in effect.

During the following discussion the differences between shared and unshared node pages were reviewed. George Penokie, summarized the clearing results on an overhead.

PRLI effect on shared mode pages: Pickup current values.

PRLI effect on unshared mode pages: Use saved values or default values if never previously logged in.

PRLO effect on shared mode pages: Clear current mode pages only if last initiator to logout.

PRLO effect on unshared mode pages: Clears current mode pages for initiator originating logout.

Bob Snively noted that this was only necessary if word 0 bit 13, establish image pair, equals 1.

A vote was taken to incorporate these changes into the clearing effects table. The result was 11 in favor, 0 opposed.

10B. FAN Timeout Questions - Dale LaFollette

Dale asked a question about target authentication with public initiators and private targets. The result of the discussion is that the target must wait RR_TOV before implicitly logging out the initiator while the host recovery must occur in E_D_TOV.

10C. Process Associators - Group

John Scheible made a presentation to the group. Behind a single host bus adapter there may be multiple processors, protocols and operating systems. Additional addressability is required to partition the data between the processes. One solution is the use of process associators, which John felt were broken. He proposed calling a meeting, which he volunteered to facilitate, and invite all of the interested parties to work to find a solution. The solution might be to fix the process associators or replace them with something else.

Dal Allan supported John's request by giving an example of one computer with a single HBA running several operating systems (LINUX, NT, MAC). There needs to be a means of providing a zoned address for each operating system. Bob Snively countered that computers already keep that sorted out. He said using a single queue into the HBA, with a thin layer of code, resolves the problem, similar to memory allocation. Dal said that this was needed to support fabric zoning, since NT is known to grab everything. There is a need to identify the routing into a switch. Each HBA currently must be in a single zone.

Aside from process associators, optional headers were proposed. Kumar Malavalli objected to their use.

An HBA can be a member of several zones, but does not provide controlled access. Roger Cummings said that an HBA can do it, but the switch needs additional information to support zoning. Bob Snively said that overlapping solutions exist. He objected to adding additional complexity.

Dal Allan suggested that a process associator could be associated with each zone. The solution must be protocol independent - not based on the TYPE field. Charles Binford proposed an alias id solution. Someone else proposed using a security header.

Reviewing the direction of the discussion, Roger Cummings corrected John's initial statement. It is not a question of insufficient addressability, but a problem of access control.

Choices include either end point or fabric control. Incorporating this in a fabric requires a trade off of cost versus value added.

In conclusion the group felt the issue should be resolved in the FC-FS group. A clear statement of problem is needed to continue the discussion.

10D. Log Select - Group. Noticed missing from FC-TAPE- The group conclusion was that this was intentional, it was reviewed and a decision was made to be silent on the issue.

11. Next Meeting Requirements: Group

12. Review New Action Items: Group

Old Action Items

#1 Dave Baldwin, Emulex. Refer default E_D_TOV issue to FC_FS. Prefer a 2 second to 10 second for point to point connections. Deferred, as Dave was not in attendance.

#2 Paul Suhler. Proposal to add Hold addition to load/unload command. Updated proposal, post to reflector and continue discussion. Deferred, as Paul was not in attendance.

#3 Dale LaFollette, write up informative annex on read errors for the FCP-2. Dale was waiting for a template from Bob Snively before completing this item.

New Action Items

#1 Dave Peterson for the next revision of the SSC. Define unique for the medium (GEM#38) and clearer statement defining sequential device positioning (GEM#39).

#2 Group, at Bob Snively's request, review clause 11, Error Recovery Procedure, in the latest FCP-2 draft and the review document, T10/99-247r1, clause 3.25, Permission to do recovery, and clause 3.28, Hold exchange information.

#3 Bob Snively, Reinstate the requirement in the FCP-2 that targets begin error recovery on the boundary required by the initiator. Modifications are required in 99-247r1 clauses 3.29, remove last paragraph, and similar errors are in clause 3.30 and 3.31. Error recovery boundary must be modulo 4 relative offset.

#4 Bob Snively to provide process associator modification notification and sample text.

#5 Bob Snively make the RSP_CODE and the SCSI Status mutually exclusive.

#6 Bob Snively. The FCP-2 shall require targets to transmit the ACC for the SRR before transmitting any recovered data. The target shall not transmit any additional non-recovered data after transmitting the ACC.

#7 Group at Bob Snively's request, review the change document clause 6.3, Incorrect use of Recovery Abort, and FCP-2 clause 11.5.1.

#8 Charles Binford Propose a new RESP_CODE 0x06: Command cleared by another initiator. Take new status code to SAM-2

#9 Bob Snively add clearing effects of PRLI on mode pages as defined in these minutes in agenda item 10A to the FCP-2.

13. Adjournment. The meeting was concluded at 5 PM

Attendance:

Dale LaFollette	STK	David Peterson	STK
Stewart Wyatt	HP	Carl Zeitler	Compaq
Dal Allan	ENDL	Edward A. Gardner	Ophidian Designs
Arlan Stone	UNISYS	Roger Cummings	DPT
Horst L. Truedtedt	True Focus	Bob Snively	Sun Microsystems
Stephen O'Neil	CMD	George Penokie	IBM
Predrag Sparic	HP	Damian Bannon	SSL
Mark Hamel	Compaq	Mark DeWilde	Pathlight
Lee Hu	Vixel	Steve Luning	Dell
John Lohmeyer	LSI Logic	Masatoshi Ichikawa	Hitachi
Kumar Malavalli	Brocade	Matt Wakeley	HP
Bill Martin	Gadzoox Networks	Gene Milligan	Seagate
Charles Binford	LSI Logic	Tim Hogleund	LSI Logic
Neil Wanamaker	Crossroads	Douglas Nast	Boeing
Ed Schurig	iphase	Masatoshi Ickikawa	Hitachi
John Scheible	IBM	Neil Baylis	Troika
Jim Nelson	Northrop Grumman		