

February 2, 2000

T10/00-139 revision 0



To: John Lohmeyer, chairperson, T10  
From: Bob Snively  
Date: February 2, 2000  
Subject: Comments on the FCP-2 letter ballot of revision 04

The following comments accompany my vote on the letter ballot about FCP-2. Those problems that are critical or key are marked with the notation “Technical \*\*\*\*\*”.

### **Sun 1 Document references**

Editorial

All sections. Document references are inconsistent or not helpful. Most people cannot hold the numbers of the various standards in their head, yet NCITS/ANSI prefers that the numbers of the referenced standards be used as the method for designating documents. Examples include:

3.4 last sentence: document cited as “FC-FS”

4.1 first paragraph: document cited as “ANSI X3.230”

6.2.7.2 first paragraph: document cited as “ANSI X3.297”

Assuming this is acceptable to the editors, I would prefer to use either the document name exclusively, or both the document name and the document number together. As examples: “FC-FS” or “FC-PH, X3.230”

### **Sun 2 Use of word FCP**

Editorial

All sections. The word “FCP” is used as a noun, in the context “The FCP...” while referring to the protocol. This looks like terrible English and reads very badly.

I would like to see the words “Fibre Channel protocol” used when speaking of the protocol, and the words “FCP standard” when referring to the document. See in particular:

Section 5.5, page 19: “The FCP” s/b “Fibre Channel Protocol”

Annex A.1, page 71. “The FCP-2” s/b “This standard”

### **Sun 3 Correct hexadecimal references**

Editorial

All sections. The conventions for hexadecimal notation (AB1Ch) are not followed consistently. Corrections need to be installed everywhere, particularly sections 6.1 (p22), 9.1 (p36), 11.3 (p60),

### **Sun 4 Update contacts**

Editorial

page ii. The E-mail addresses for the X3T10 chair and the T10 reflector must be updated. The SCSI BBC information must be updated. The references to X3 must be changed to NCITS.

## **Sun 5 Remove document revision history**

Editorial

page iii. The document revision history should be removed.

## **Sun 6 Combine annexes into primary table of contents**

Editorial

page viii. The annex table of contents should be moved from page x and appended to the normal table of contents on page viii.

## **Sun 7 Correct document description**

Editorial

Section "Introduction", page xv.

Clause 7 should be inserted in the document description with text that says: "Clause 7 describes the FC-4 specific name server object for FCP."

Annex E should be inserted with text that says "Annex E is an informative annex providing examples of error recovery procedures."

Annex G should be inserted with text that says "Annex G is an informative annex showing examples of ELS formats required for proper FCP-2 recovery operations."

The text describing removal of annexes should clarify that the referenced document is the old standard.

Text for annexes H, I, and J should be inserted.

## **Sun 8 Remove redundant sentence**

Editorial

Page 2, Clause 2.3, first paragraph, last sentence is redundant and should be deleted.

## **Sun 9 Clarify definition of base address**

Editorial

Page 2, clause 3.1.5, the definition should be replaced with "base address: The virtual address of the byte having the lowest address among the bytes to be transferred to or from an application client buffer."

## **Sun 10 Data overlay definition restriction**

Editorial

Page 3, clause 3.1.13, the definition of data overlay should be corrected to exclude link recovery. "data overlay: Data overlay occurs when data is transferred to or from the same offset of the SCSI application client buffer more than once during the same command, except for the recovery of link transmission failures."

## **Sun 11 Circular definition of MCM**

Editorial

Page 3, 3.1.27, the definition of MCM should have the word "MCM" removed in the defining text in two locations. A reference to FC-AL-3 should be provided.

## **Sun 12 typo**

Editorial

Page 5, section 3.3.8. “standards” s/b “standard”.

## **Sun 13 typo**

Editorial

Page 6, section 3.4. second line s/b “...defined in the glossary or in the text...”

## **Sun 14 Clarify number of sequences**

Editorial

Page 8, section 4.1, next to the last sentence s/b “The maximum number of active sequences that can simultaneously be open between an initiator FCP\_Port and a target FCP\_Port is restricted by the allowable range of values of the Sequence ID to 256, as defined in FC-PH.”

## **Sun 15 Use of “the FCP”**

Editorial

Page 8, section 4.2, first paragraph. The word “FCP” in this sentence should be replaced with “FCP host adapter” in three places.

## **Sun 16 typo**

Editorial

Page 8, section 4.2, 4th paragraph. “... command, has...” s/b “...command, and has...”

## **Sun 17 typo**

Page 8, Correct fonts in section 4.2, 5th paragraph.

## **Sun 18 table typo**

Page 11, section 4.6, Remove extra line in table 2.

## **Sun 19 Clarify task management completion**

Editorial

Page 12, section 4.7, second paragraph. Change “A task management function ends with an FCP\_RSP IU that indicates whether it was correctly accepted.” to “A task management function ends with an FCP\_RSP IU that indicates the completion status of the function.”

## **Sun 20 Clearing effects of PRLI/PRLO**

Technical

Page 13, table 4. In the column that indicates the clearing action for PRLI/PRLO, all the “Y” entries actually only apply for the affected image pair. This must either be indicated in a note or have a separate entry for that case in each relevant row.

## **Sun 21 Clearing effects on buffered data**

Technical

This was presented by George Penokie on Jan 26, 2000. He suggests that reserved XOR data be cleared if there is a target power cycle, a reset LIP, a logout of all initiators, TPRLO, SCSI target reset, or a Logical Unit Reset. All other cases would preserve the data.

## **Sun 22 Clarification of mode page management**

Editorial

Page 14, table 5, column 4. The column should indicate that this is the state after the PRLI/PRLO has been executed.

## **Sun 23 typo**

Editorial

Page 15, section 4.9, title. The title should indicate that this references only Process Login/Logout.

## **Sun 24 Process Login image definition**

Technical

Page 16, section 5.1, second paragraph. In this paragraph, it is clearly stated that the process associator does not take place in the identification of the initiator or target. However, in the third sentence, there is a left-over sentence that indicates that more than one logical initiator or logical target image may be defined by the process associator. The sentence "More than one logical initiator or logical target image may be defined..." should be deleted.

## **Sun 25 Process Associator for FCP\_Port addressability**

Technical

Page 16, section 5.2. This section defines a process associator value for third-party referencing of FCP addresses. Since the Process Associator does not take part in the initiator or target definition, it is not necessary to include the Process Associator in the definition. Table 7 should have the PA\_VAL bit changed to reserved and the Process Associator field changed to reserved. Section 5.2.1 and 5.2.3 should be deleted.

## **Sun 26 Incorrect definition of Data Out IU T7**

Editorial

Page 18, section 5.4, table 8. This is an editorial error that was not caught in the original FCP document. The T7 Data Out action IU can only occur when there are two consecutive write data sequences. With the final definition of Disable Write Data Transfer, all FCP\_DATA IUs are separated by a XFER\_RDY IU, making T7 an unused IU. T7 should be removed from the table and notes of table 8.

## **Sun 27 Remnant of I2 IU needs to be removed**

Editorial

Page 19, table 9, notes. The I2 data IU has been removed from the table. The third note should also have it removed.

## **Sun 28 Clarify definition of tag**

Technical?

Page 21, section 5.6.9. The third sentence indicates that the OX\_ID is the tag defined by SAM. That is not strictly true, because the OX\_ID exists for all exchanges, even those that have no tag definition. The sentence should be. "The value of the OX\_ID is used to identify an FCP I/O Operation the same way that the tag value identifies I/O Operations in ANSI X3.270."

## **Sun 29 Correct RO requirements**

**Technical \*\*\*\*\***

Page 21, section 5.6.11, last sentence. The sentence indicates that RLTV\_OFF is not required if both FCP\_Ports can unambiguously reassemble the transmitted IUs. This creates severe interoperability problems for those ports that may be attached that cannot perform this magic. This interoperability issue is not negotiated in any login parameters. The correct solution is to require the presence of the RO, then allow the recipient of the data to use the RO or other FC-PH mechanisms to reassemble the data.

## **Sun 30 Make Process Associator obsolete for FCP**

**Technical \*\*\*\*\***

Process Associators create a complex functionality that cannot successfully distinguish separate images in the initiator or the target. In the target, separate images (including both initiator and target enforced protection) are created using the logical unit. In the initiator, no separate images are explicitly defined, but they can be emulated by using more than one port address identifier for the initiator port. Since the function originally conceived of (but never practically implemented) by Process Associators in FCP can be done practically by other mechanisms more natural to both FC and SCSI, Process Associators should be made obsolete for FCP. Process Login should still be used to negotiate capabilities and to identify target/initiator pairs.

## **Sun 31 Distinguish image pair and initiator/target pair**

**Technical \*\*\*\*\***

Section 6, all sections. In many places, image pair is referenced. However in some cases it is implicitly a relationship between initiator and target, and in other cases it is a relationship between initiator process image and target process image. The two cases should be distinguished by using the word "image pair" for those that use a process image and "initiator/target pair" for those that do not use a process image.

## **Sun 32 typo**

Editorial

Page 22, section 6, first paragraph. "extended link services in ANSI" s/b "extended link services defined in ANSI"

## **Sun 33 Correct login requirement**

Editorial

Page 22, section 6, second paragraph, last sentence: is "Devices introduced into a configuration or modifications in the addressing or routing of the configuration may require new login procedures.", should

be “Devices introduced into a configuration or modifications in the addressing or routing of the configuration may require the login and discovery procedures to be executed again.”

### **Sun 34 Problem with process associators**

Technical

Page 23, section 6.2.2: Refer to Sun 30. The originator for all FCP communications is the initiator. The mechanisms to manage multiple images behind a single initiator port are incomplete. Section 6.2, third paragraph is one example of this attempt to create an unsupported function. Similar problems exist in the corresponding paragraph of 6.2.1 and 6.2.3.

### **Sun 35 Problem with clearing PRLI image pairs**

Editorial

Page 24, section 6.2.5, first paragraph. The third sentence of this paragraph indicates how outstanding exchanges are affected by a PRLI. This is actually referenced in table 4, but incorrectly. See Sun 20. The correct solution is eliminate Process Associators. Failing that, a reference should be used here and the definition of the clearing effects should be placed in table 4.

### **Sun 36 Correct behavior of new PRLI**

Editorial

Page 24, second paragraph, reads in part:

Immediately after the execution of the first PRLI, both members of all image pairs shall have the same state as they would have after a hard reset or a power on with respect to each other. No tasks, reservations or status shall be present in either SCSI device. The MODE SELECT parameters will assume their default or saved states for all image pair. Tasks, reservations, status, and MODE SELECT parameters for other initiators are not affected. A Unit Attention condition (Sense Key = 6) with an Additional Sense Code of Reset Occurred (ASC = 29, ASCQ = 00) shall be presented upon the first attempt to communicate between the N\_Ports using FCP when a new PRLI has been performed. A target port shall not generate a unit attention condition for initiators which are already logged in. Subsequent PRLI operations shall have no effect on FCP operation between two devices except where new requirements are negotiated between the devices.

The text should be corrected as follows:

Immediately after the execution of ~~the first any~~ PRLI, both members of all ~~new~~ image pairs shall have the same state as they would have after a hard reset or a power on with respect to each other. No tasks, ~~non-persistent~~ reservations or status shall be present in either SCSI device. The MODE SELECT parameters will assume their default or saved states for ~~all the new~~ image pairs. Tasks, reservations, status, and MODE SELECT parameters for other ~~initiators image pairs~~ are not affected. A Unit Attention condition (Sense Key = 6) with an Additional Sense Code of Reset Occurred (ASC = 29, ASCQ = 00) shall be presented upon the first attempt to communicate between the N\_Ports using FCP when a new PRLI ~~image pair~~ has been ~~performed created~~. A target port shall not generate a unit attention condition for ~~initiators initiator members of image pairs~~ which are already logged in. Subsequent PRLI operations shall have no effect on FCP operation between two devices except where new requirements are negotiated between the devices.

This is another example of the problems associated with Process Associators.

### Sun 37 Correction to PRLI request

Technical

Page 25, section 6.2.6.4. If process associators are removed for FCP, the value for the Establish Image Pair field shall be 0.

### Sun 38 Definition of process suspect

Editorial

Page 26, section 6.2.6.7, first paragraph. The word “process” in the first sentence should probably be replaced with the words “member of the image pair”, or alternatively, the words initiator and target should be used.

### Sun 39 Incorrect use of word “image pair”

Editorial

Page 26, section 6.2.6.7, third paragraph. The paragraph presently reads:

An image pair shall use the retransmission capability only if the RETRY bit is set in both the request payload and in the accept payload. If the RETRY bit is set to 0 in either the request payload or the accept payload, the SRR shall not be performed by the initiator. If the SRR is received by the target, the SRR shall be rejected with LS\_RJT.

The text should read:

An ~~image pair~~ initiator and target port pair shall use the retransmission capability only if the RETRY bit is set in both the request payload and in the accept payload. If the RETRY bit is set to 0 in either the request payload or the accept payload, the SRR shall not be performed by the initiator. If ~~the an~~ SRR ELS is received by ~~the a~~ target that has set the RETRY bit to 0, the SRR shall be rejected with LS\_RJT.

### Sun 40 Correct Write XFER\_RDY Disabled definition

Technical

Page 27, section 6.2.6.13, first two sentences should be rewritten as follows:

When this bit is set to 0, FCP\_XFER\_RDY IUs shall ~~be used~~ transmitted by the target to request each of the SCSI write FCP\_DATA IUs from the initiator. for SCSI write operations. When this bit is set to 1, FCP\_XFER\_RDY IUs ~~may be not~~ shall not be used before the first FCP\_DATA IU to be transferred in the write operation.

### Sun 41 Complete Image Pair Established definition

Technical

Page 28, section 6.2.7.1, first paragraph should be rewritten as follows:

IMAGE PAIR ESTABLISHED is valid only if bit 13 was set to 1 on the corresponding Service Parameter page of the PRLI request and if the image pair was correctly established.

### Sun 42 Correct PRLO text

Technical \*\*\*\*\*

Page 29, section 6.3, first paragraph, should have the same corrections applied as those defined in Sun 36. In addition, it should be clarified that tasks are reset for all image pairs that have been “unpaired” and destroyed by the PRLO, but not for other image pairs.

### **Sun 43 State after PRLO**

Technical

Page 29, section 6.3, paragraph 3, second sentence. The second sentence should be modified to read:

After PRLO, no further FCP communication is possible between those two N\_Ports.

### **Sun 44 State of image pairs after PRLO**

Technical

Page 29, section 6.3, 4th paragraph. The paragraph leaves some uncertainty about the proper response to the PRLO and the proper state of image pairs if some are discontinued, some are not discontinued, and some do not exist (or never existed). I believe that the PRLO should respond as if the image pairs that do not exist are successfully discontinued just as if they existed.

### **Sun 45 Clarify table 13**

Editorial

Page 30, section 7.1, table 13. The table should be restructured in bit/byte format to make the definitions clearer. The fields should be identified and described by field name.

### **Sun 46 Verify FCP specific object format**

Editorial

Page 30, section 7.2. This section was written before FC-GS-3 was available. It should be examined for consistency with FC-GS-3 and any corrections installed.

### **Sun 47 Install references**

Editorial

Page 31, section 8, table 15. The references for FCP\_ACC and FC\_\_RJT need to be installed.

### **Sun 48 Clarify table 16**

Editorial

Page 32, section 8.1, table 16. The table should be restructured in bit/byte format to make the definitions clearer. The fields should be identified and described by field name.

### **Sun 49 SRR inconsistent error reporting**

Technical

Page 31, section 8.1. The seventh paragraph indicates that an SRR that cannot be accepted is treated as an “Initiator Detected Error.” The third paragraph of the “payload” description on the next page indicates that such an error will be indicated with an FCP\_RJT. These two statements need to be reconciled. I expect that there are really two cases. One could be treated as an Initiator Detected Error (failed link recovery), while the other could be treated as an FCP\_RJT (invalid payload contents).



## Sun 50 Describe payload for FCP\_RJT

Editorial

On page 33, section 8.3, the paragraphs describing payload, the description is complicated enough to warrant the creation of a table that describes the complete payload.

## Sun 51 FCP\_LUN format

Technical

On page 36, section 9.1.1.1. FCP-2 revision 4 shows the FCP\_LUN field in the FCP\_CMND IU as an 8 byte field. It states: "The FCP logical unit number (FCP\_LUN) is the address of the desired logical unit in the attached subsystem. The FCP\_LUN field is specified by ANSI X3.230 for all IUs of Category 6."

The last sentence implies one can find detail on the contents of the LUN field in X3.230, but that document is FC-PH, and FC-PH simply defines the field as "Entity Address (FC-4 dependent)". So the references point at each other.

Was the intent in FCP-2 to let the LUN field be defined by the appropriate SCSI standard document(s)? FCP mentioned the SCSI Device Model. Or was more specific guidance to a specific format of the LUN field intended in FCP-2? FCP Annex C gave a SCSI Controller Command LUN field usage example.

The proper documentation is probably in SAM-2.

## Sun 52 FCP\_CMND IU sub-section titles

Editorial

The titles for sections 9.1.1.x starting on page 37 should not have the byte number included in the title.

## Sun 53 Clarification of ordered queueing

Technical

Page 37, section 9.1.1.3, the text of paragraph 4 should be rewritten as follows:

ORDERED\_Q requests that the task be managed according to the rules for an ORDERED task attribute. With a class 2 fabric, special care ~~must be~~ should be taken exercised to guarantee successful ordering. ~~Sequential In order~~ delivery ~~must should~~ be requested at login to ensure correct ordering among tasks. ~~FCP\_CMND IUs must be acknowledged before new FCP\_CMND IUs are issued to avoid inadvertent reordering of commands during retries of F\_BSY. Acknowledgements should be received before new FCP\_CMND IUs are issued to avoid inadvertent reordering of commands during delays in the fabric, including retries of F\_BSY.~~ Ordering can also be accomplished by waiting for the completion of those commands requiring ordering before transmitting the FCP\_CMND for the next FCP I/O operation, or by using the precise delivery mechanism.

## Sun 54 Clarification of task management flags

Editorial

Page 38, section 9.1.1.4, the words "Task Management function" s/b "Task Management request".

## Sun 55 ACA clarification

Technical

Page 38, section 9.1.1.4, Clear ACA. It is unclear what the proper behavior is if there is no ACA present when a Clear ACA is transmitted. References to SAM should be provided to clarify this.

## **Sun 56 Complete Clear ACA description**

Technical

Page 38, section 9.1.1.4, Clear ACA, 4th paragraph, the incomplete sentence should be completed as:

Depending on the MODE SELECT parameters that have been established, additional FCP I/O operations may have to be aborted by the recovery abort as part of the process of clearing the automatic contingent allegiance.

## **Sun 57 Clarify Target Reset**

Editorial

Page 38, section 9.1.1.4, Target Reset. The first paragraph should be rewritten as a list (see Logical Unit Reset text as an example) to make it easier to read.

## **Sun 58 Correction of note about SAM-2**

Editorial

On page 39 and 40, section 9.1.1.4, the text of the notes about SAM-2 should be corrected by changing the words “by this mechanism” to “with this completion status”. There are 4 such notes to be corrected.

## **Sun 59 Clarify dependant logical unit**

Editorial

Page 39, section 9.1.1.4, Logical Unit Reset, item 6. The definition of dependent logical units is a little vague. It is not included in the glossary and the reference 4.11 does not exist. The definition should be extracted from SAM-2, placed in the glossary, and referenced here.

## **Sun 60 Clarify logical unit reset**

Editorial

Page 39, 9.1.1.4, logical unit reset, second paragraph after list. The first sentence should be re-ordered to say: “shall be terminated using a recovery abort by whichever port”.

## **Sun 61 Clarify logical unit reset**

Editorial

Page 39, 9.1.1.4, logical unit reset, last paragraph. The Logical Unit Reset does not address targets, but rather logical units. The ambiguity should be addressed in terms of those exchanges ambiguous with respect to the logical unit, not those ambiguous with respect to the target.

## **Sun 62 Clarify Clear Task Set**

Editorial

Page 40, 9.1.1.4, Clear Task Set. This section has the same problems as the logical unit reset, except that the focus of the logical unit reset should be logical unit and the focus of clear task set should be task sets. See Sun 60 and Sun 61.

## **Sun 63 Clarify Additional FCP\_CDB length**

Technical

Page 40, 9.1.1.5. Add a sentence to this section. “The Additional FCP\_CDB Length field shall be zero for task management requests.”

### **Sun 64 Clarify Additional FCP\_CDB**

Technical

Page 41, 9.1.1.9. The text of the second sentence should be changed as shown: “The ADDITIONAL\_FCP\_CDB ~~is not valid and is ignored~~ **shall not be present** if any task management flag is set to 1.

### **Sun 65 Use of word FC-PH**

Editorial

Page 41, 9.1.2.2, second paragraph. The text “The FC-PH allows...” s/b “The FC-PH standard allows...”

### **Sun 66 Verify proper execution of recovery abort**

Technical?

Page 41, 9.1.2.2, fourth paragraph. The text “A target ... specified RX\_ID.” needs to be clarified. In particular, it is not clear what the detailed values of the recovery qualifier are with respect to RX\_ID.

### **Sun 67 Clarify FCP\_XFER\_RDY**

Editorial

Page 42, section 9.2, first paragraph. The text “... to perform ...” s/b “... to receive ...”.

### **Sun 68 Clarify FCP\_XFER\_RDY when disabled**

Editorial

Page 42, section 9.2, second paragraph. An additional sentence should be added at the end of the paragraph as follows: “The first FCP\_DATA IU is transmitted without a preceding FCP\_XFER\_RDY.”

### **Sun 69 Clarify obligation of initiator**

Editorial

Page 42, section 9.2, third paragraph. The last sentence should be changed to read: “The initiator shall be ready to transmit ~~the entire~~ **any part or all of the** FCP\_DL bytes of data.”

### **Sun 70 Clarify Data\_RO**

Editorial

Page 42, 9.2.1. The first sentence should change the words “the next FCP\_DATA” to “the requested FCP\_DATA”. In addition, a reference should be put in place for the SCSI-3 application client buffer offset, probably in SAM-2.

### **Sun 71 Clarify Burst\_LEN**

Editorial

Page 42, 9.2.2, first paragraph should be changed as follows: “~~For data transfers from the SCSI initiator to the target,~~ The BURST\_LEN field indicates the amount of buffer space prepared for the next

FCP\_DATA IU and requests the transfer of an IU from the initiator of that exact length. This value is the same as the SCSI data delivery request byte count. See SAM-2.

### **Sun 72 Clarify maximum burst length**

Technical

Page 42, 9.2.2, third paragraph. The following text should be added to the third paragraph: “A BURST\_LEN greater than FCP\_DL or longer than the maximum burst length specified by the disconnect/reconnect mode page is not valid.”

### **Sun 73 Clarify FCP\_DATA IU**

Editorial

Page 43, 9.3, first paragraph. The first paragraph should be changed as follows: “The data associated with a particular FCP I/O Operation is transmitted in the same exchange that sent the FCP\_CMND requesting the transfer. identified by the FQXID.”

### **Sun 74 Simplify Mode Sense/Select reference**

Editorial

Page 43, section 9.3, second paragraph. Delete the last sentence and replace it with a reference to 10.1.1.6.

### **Sun 75 Improve text**

Editorial

Page 43, section 9.3, fourth paragraph. Since there is only one Data Out IU, the parenthetic (T6 or T7) should be deleted. The last sentence needs to be corrected to indicate “first” rather than “corresponding”.

### **Sun 76 Improve text**

Editorial

Page 43, section 9.3, fifth paragraph. Since there is only one Data Out IU, the parenthetic T6 and T7 should be removed. The last sentence needs to be corrected to read: “The command is completed normally except that data beyond the FCP\_DL count shall not be transferred and that the appropriate overrun condition is presented. for presentation of the overrun condition. See 9.4.1.”

### **Sun 77 Clarify data is contiguous**

Editorial

Page 43, section 9.3, 8th paragraph. The third sentence “The target shall not request that sets of data in the middle of a transfer not be transferred.” should be deleted. The second sentence covers this.

### **Sun 78 Stylistic improvement**

Editorial

Page 44, section 9.3, last paragraph. The wording of the following sentence, “ANSI X3.230 specifies the mechanisms by which an IU shall be transferred.” should be changed to “ANSI X3.230 specifies how an IU shall be transferred.”

## Sun 79 Improve description of linking

Editorial

Page 44, section 9.4, third paragraph. The wording of the last sentence should be changed to: “~~There is no FCP-2 function equivalent to the~~ The LINKED COMMAND COMPLETE or LINKED COMMAND COMPLETE (WITH FLAG) function defined by SAM and SAM-2 is implicit in the presentation of the proper status in the FCP\_RSP.”

## Sun 80 Include task management in FCP\_RSP\_INFO

Technical

Page 47, section 9.4.10. The FCP\_RSP\_INFO description should contain an explicit requirement that FCP\_RSP\_INFO is always present in a task management response.

## Sun 81 Improve table format

Editorial

Page 48, section 9.4.10, table 27. The table should be reformatted to clearly define the bits and bytes used.

## Sun 82 Correct task management completion

Technical

Page 48, section 9.4.10, last paragraph. The first sentence should be changed as follows: ~~The task management function may or may not have been performed by the target if RSP\_CODE is returned or if no FCP\_RSP is returned before the Exchange is aborted.~~ The completion status of the task management function is indicated by the RSP\_CODE. If the Exchange is aborted before the FCP\_RSP is returned, the completion status is unknown.

## Sun 83 Capitalize error codes

Editorial

Page 48, section 9.4.10, table 28. SPI-3 has elected to place the packetized failure codes in upper case. They recommend that the corresponding response codes of table 28 all be upper case.

## Sun 84 Verify task management completion

Technical

There was a statement at one meeting that the FCP-2 document is not consistent with the SAM-2 document with respect to the task management function completion codes. This must be verified.

## Sun 85 Correct description of SCSI mode parameters

Editorial

Page 50, section 10.1, first sentence: The sentence should be changed to read: “This clause describes the block descriptors and the pages used with MODE SELECT and MODE SENSE commands that ~~influence~~, control and report the parameters that influence the behavior of FCP.”

## Sun 86 Clarify requirements for parameters

Editorial

Page 50, section 10.1.1, last sentence: The sentence should be changed to read: “If a parameter that is not appropriate for the an standard for FCP-2 SCSI-3 device is set nonzero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

### **Sun 87 Add recommendation to Bus Inactivity Limit**

Technical

Page 52, section 10.1.1.3: The following note should be added after the last paragraph:

Note: Because of the low overheads associated with initiating and closing bus tenancy on Fibre Channel links, device servers should end tenancies immediately upon completing the required transfers.

### **Sun 88 Note that FC Port Control page violates standards**

Editorial

Page 54/55, section 10.1.3: The following note should be added before Table 32:

Note: Some of the bits defined by the Fibre Channel Port Control page require the port to violate one or more of the fibre channel standards. The non-standard behaviors have been identified as useful for certain specialized operating environments.

### **Sun 89 typo**

Editorial

Page 55, section 10.1.3.2:

Initiated s/b Initiated  
sequences s/b sequence

### **Sun 90 Correct DTIPE bit = 0 description**

Technical

Page 55/56, section 10.1.3.2: At present, if DTIPE is set to zero, vendor specific initialization is expected. The proper behavior should be to follow the initialization method specified by FC-AL-2.

### **Sun 91 Improve RHA readability**

Editorial

Page 56, section 10.1.3.4: The second paragraph should be divided into two paragraphs separated between “... get its hard address.” and “If the hard address ...”.

### **Sun 92 Clarify RR\_TOV default**

Technical

Page 57, section 10.1.3.9: The next to the last sentence should be changed to read: If no timer is specified, the RR\_TOV value in byte 7 shall be ignored by the device server and a vendor specific default value shall be used.

### **Sun 93 typo**

Editorial

Page 60, section 11.2. "... with the aborted Sequence." s/b "... with an aborted Sequence."

#### **Sun 94 Concern about organization of error recovery section**

Editorial

Page 62, section 12. It appears that a descriptive paragraph or model would be appropriate either here or in section 4.0. The overall structure of 12 depends on two types of error detection, one that works for all classes of service and an additional one that works for acknowledged classes of service. Once an error is detected, there appears to be two types of error recovery that can be performed, one that does exchange level recovery, and the other that does sequence level recovery. This is a bit difficult to pick out of the document.

#### **Sun 95 Remove redundant sentence**

Editorial

Page 62, section 12.1.1. The sentence "An FCP-2 ... defined below." should be deleted. It is left over from a previous revision of the document.

#### **Sun 96 Emphasize optional error recovery**

Editorial

Page 62, section 12.1.2. The first sentence should be rewritten as: "SCSI devices may use the mechanisms described in this chapter to detect the presence of link errors, then perform **optional** retransmission procedures that **will** allow the commands to be completed without requiring **complex** higher level recovery algorithms." The extra line space above the paragraph should be deleted.

#### **Sun 97 typo**

Editorial

Page 62, 12.1.2, third paragraph. Correct font.

#### **Sun 98 Clarify error detection**

Editorial

Page 62, section 12.2.1, first paragraph. The sentence should be rewritten as: "The Exchange originator (SCSI Initiator) **may detect the following errors. It may optionally further identify and recover the error as described in 12.3.** ~~shall initiate error detection and recovery described in 12.3 for the following:~~"

In addition, in line item 3, "an Sequence" s/b "a Sequence".

The same rewrite should be done for the paragraph associated with the Exchange responder.

#### **Sun 99 Clarify sequence error detection**

Editorial

Page 63, sections 12.2.1 and 12.2.2. Item 4 of the target list of section 12.2.2 really applies to all classes and should be deleted from 12.2.2. In 12.2.1, the wording for sequence errors should be changed to match the text deleted from 12.2.2. Reference to section 12.3.9 for the recovery process should be made.

## Sun 100 Clarify error detection

Editorial

Page 62, section 12.2.2, first paragraph. The second sentence should be rewritten as: “The Exchange originator (SCSI Initiator) may detect the following errors. It may optionally further identify and recover the error as described in 12.3. ~~shall initiate error detection and recovery described in 12.3 for the following:~~”

The same rewrite should be done for the paragraph associated with the Exchange responder.

## Sun 101 Clarify exchange level error recovery

Editorial \*\*\*\*\*

Section 12.1.1 outlines how exchange level error recovery works. The only place it is actually described in detail is in 12.3.8 , 12.3.9, 12.4, and 12.5 actually describe different portions of the exchange level error recovery. However, 12.3.8 and 12.3.9 in large measure duplicate the detection discussion of 12.2.1 and 12.2.2. Section 12.4 overlaps with the recovery information in 12.5.2. I believe that 12.3.8 and 12.3.9 should be selectively deleted where the information is duplicated by 12.2.1 and 12.2.2. I believe that section 12.4 should be carried into section 12.5.2, with which it is almost totally redundant.

## Sun 102 Clarify exchange level error recovery

Editorial

Page 63, Section 12.2.2 The last two paragraphs describe recovery mechanisms, not detection mechanisms. These belong in another section, possibly 12.5.

## Sun 103 Overall formatting of recovery suggestion

Editorial

The relationship among recovery algorithms and detection algorithms is not as clear as it should be. I would suggest the following organization:

12.1 Overview

12.1.1 Overview of exchange level recovery

12.1.2 Overview of sequence level recovery

12.2 Initial FCP error detection

12.2.1 Error detection for all classes of service

Initiator

Target

12.2.2 Additional error detection for acknowledged classes

Initiator

Target

12.3 Exchange level error recovery (largely the same as old 12.5)

12.3.1 SCSI initiator abort of exchange (largely the same as old 12.5.1)

12.3.2 SCSI target abort of exchange (largely the same as old 12.5.2, combined with 12.4)

12.4 FCP-2 specific error recovery (this is distinguished by special use of REC and time-outs.)

This contains all the sections from 12.3.1 to 12.3.7. Note that 12.3.8 and 12.3.9 are included already in 12.2.1.



12.5 Second level error recovery (This contains all the sections from 12.6.1 through 12.6.3)

12.6 Responses to FCP-level frames before PLOGI or PRLI (This contains all of 12.7)

### **Sun 104 Clarify REC polling**

Editorial

Page 64, section 12.3.1. The overall model of polling using REC is never defined. Parts of it are included in the REC\_TOV definition, some implicit timeout conditions in 12.2.1, and parts in 12.3.1. Since we are not covering the overall error recovery in section 4, it seems appropriate to spend a paragraph indicating how polling is performed in 12.3.1.

### **Sun 105 Clarify REC response if no OX\_ID**

Technical

Page 64, section 12.3.2. The definition of the reason code for the LS\_RJT is incomplete. The definition should be: "... for the REC with a reason code of Logical Error and a reason code explanation of Invalid OX\_ID-RX\_ID combination (0317h)."

### **Sun 106 Correct error recovery reason**

Editorial

Page 64, section 12.3.2. The sentence "This is to ensure that no reply Sequences have been lost." is really not correct. It will eventually become apparent that they were lost. This is really to find out more quickly that the sequences have been lost.

### **Sun 107 Clarify REC response if no OX\_ID**

Technical

Page 66, section 12.3.7. The definition of the reason code for the LS\_RJT is incomplete. The definition should be: "... to the REC from the target will be a LS\_RJT with a reason code of Logical Error and a reason code explanation of Invalid OX\_ID-RX\_ID combination (0317h)."

### **Sun 108 Exchange bashing options**

Technical

Page 66, section 12.3.8. The initiator may also abort the exchange with any task management function or with the ABORT TASK function, which uses the recovery abort protocol, which uses ABTS. It is my impression that ULP\_TOV will use one of the higher level functions, probably ABORT TASK, to invoke the ABTS.

### **Sun 109 Redundant sections**

Technical

Page 67, section 12.4 appears to be largely redundant with 12.5.2 and should be combined with it.

### **Sun 110 Exchange bashing options again**

Editorial

Page 67, 12.5.1, paragraphs 4, 5, and 7. The words "ABTS protocol" should be replaced with "recovery abort"

## Sun 111 Stylistic correction

Editorial

Page 69, section 12.6.1. The phrase “If the SCSI target is not on a remote loop,” s/b “If the SCSI target is on the local loop or if the loop is private,”.

## Sun 112 Clear resources after second level error recovery of REC

Technical

Page 69, section 12.6.2, next to last paragraph. The paragraph should be modified to read: “The REC shall be retried at a rate not to exceed once per the timeout period for at least 3 times. If none of the RECs receive a response, the Initiator shall report an error condition to the ULP, clear resources associated with the exchange, and perform an implicit logout with the target.”

## Sun 113 Clear resources after second level error recovery of SRR

Technical

Page 69, section 12.6.2, next to last paragraph. The paragraph should be modified to read: “The SRR shall be retried at a rate not to exceed once per the timeout period for at least 3 times. If none of the SRRs receive a response, the Initiator shall report an error condition to the ULP, clear resources associated with the exchange, and perform an implicit logout with the target.”

## Sun 114 Clarify sending of logout

Technical

Page 69, section 12.7. The first paragraph should be rewritten as follows: “If a SCSI Target receives an FCP\_CMND from an NL\_Port a port with which it has not successfully completed N\_Port Login (PLOGI), it shall discard the FCP\_CMND and, in a new exchange, send LOGO to that NL\_Port port. No Exchange is created in the SCSI Target for the discarded request, and the originator of the discarded request terminates the Exchange associated with the discarded request and any other open Exchanges for the SCSI Target sending the LOGO. ~~The LOGO is not part of the Exchange associated with the discarded request.~~”

## Sun 115 Allow implicit login

Technical

Page 69, section 12.7. The successful completion of a login should include an implicit login. A new paragraph should be added at the end that says: “FCP-2 devices that have used implicit PLOGI and/or implicit PRLI to establish their parameters and relationships may accept all FCP-2 IUs exactly as if they had completed an explicit PLOGI and/or PRLI.”

## Sun 116 Remove placeholder

Technical

Page 69, section 12.7. Delete the last sentence of the section, which had been reserved as a placeholder for any other frames of interest.

## Sun 117 Remove editor’s note

Editorial

Page 71, section A.1. The editor’s note should be removed.

## **Sun 118 Remove SCSI Parallel Interface**

Editorial

Beginning page 74, sections of A.5. These sections were originally taken from a parallel SCSI document. The words “SCSI parallel interface services” s/b “Fibre Channel Protocol for SCSI services”.

## **Sun 119 Resetting FCP**

Editorial

Page 75, section A.5.1.1. The words “ABORT TASK message” s/b “ABORT TASK function”.

## **Sun 120 Create change document for FC-FS**

Editorial

Page 77, section B.1. The words “a future version of FC-PH” should be “FC-FS”.

The document requesting these changes for FC-FS must be prepared by the editor.

## **Sun 121 Clarify Basic Link Services requirements**

Editorial

Page 77, section B.2. The word “exceptions” s/b “additional functions”.

## **Sun 122 Clarify ABTS description**

Editorial

Page 77, section B.2.1, second paragraph should be rewritten to say: “~~FC-PH does not specify a mechanism to determine which behavior an ABTS should create. While using~~ **To meet the requirements of** the FCP-2 ~~standard protocol~~, the default value of bit 0 in the ABTS request parameter field shall be interpreted as requiring the aborting of the exchange, as described in section 9.1.2.2 on page 41 of this standard. ~~While using the FCP-2 protocol, a~~ **A** value of 1 in bit 0 of the parameter field requires that the sequence be aborted as described in ~~FC-FS FC-PH, clause 21.2.2.1 and as described in section 12.4 on page 67 of this standard.~~

## **Sun 123 Reference update**

Editorial

Page 77, section B.3. “FC-PH” s/b “FC-FS”.

## **Sun 124 Clarify REC description**

Editorial

Page 78, section B.3.1, first paragraph. The text “If the RX\_ID is unspecified in the request” s/b “If the RX\_ID is specified as undetermined in the request”.

## **Sun 125 Reference update**

Editorial

Page 78, section b.3.1, “FC-PH” s/b “FC-FS”.

## **Sun 126 Clarify REC Accept payload**

Editorial

Page 79, section b.3.1, table B.4 and text underneath:

- a) Table B.4 should be modified to show the byte/bit layouts
- b) The first sentence under the table, “E\_STAT ...” under the table should be deleted. The full byte definitions of the E\_STAT value should be incorporated in table B.4.
- c) The third paragraph under the table should be changed from “set the” to “set to the”.

## **Sun 127 typo**

Editorial

Page 81, section C.1: “initiators and targets” s/b “initiators or targets”.

## **Sun 128 document format improvement**

Editorial

Page 84, C.2, The text should be moved to join Figure C.1.

## **Sun 129 Consider T11/99-722v2**

Technical

Carl Zeitler of Compaq has offered document T11/99-722v2 for consideration with respect to error recovery examples including those in Annex D. If these considerations are not included in his formal comments, they are included in this formal comment.

## **Sun 130 Correct error recovery procedure**

Technical

Page 95 and 96, Figure D.7 and D8. The last sentence in figure D.7 should be rewritten to read: “The Target retransmits the FCP\_XFER\_RDY using the specified Relative Offset ~~(or a Relative Offset smaller than the Relative Offset specified in the SRR in order to be aligned on an appropriate boundary in the Target).~~”

In addition, the label for the last data transfer arrow should be: “FCP\_DATA (seq=2, cnt=1)”.

## **Sun 131 typo**

Editorial

Page 97 and 98, Figure D.9 and D.10. The label for the last data transfer arrow should be “FCP\_DATA (seq=2, cnt=1)”.

## **Sun 132 typo**

Editorial

Page 100, Figure D.12. The last two sentences need to be separated by a blank space.

## **Sun 133 Acknowledged classes**

Editorial

Page 103, Table E.1. The words “Class 2 or Class 3 Frame” s/b “Acknowledged or unacknowledged frame”. The words “Class 2 only frame” s/b “Acknowledgement frame”.

### **Sun 134 Clarify text of E.2**

Editorial

Page 103, section E.2 (all). The example weaves together queued and unqueued cases. As a result, it is harder than necessary to interpret these pages. The section should separate the queued and unqueued cases into two separate examples.

### **Sun 135 Implicit confirm?**

Technical

Page 104, Section E.2.2. The third sentence of the second paragraph is not correct. There is no concept of an implicit confirmation with respect to a target-initiator nexus in SCSI or FCP. The best solution is probably to delete the sentence.

### **Sun 136 Is example desirable?**

Editorial

Page 105, Figure E.1. After some review, this picture looks just like D.9 and D.10. If this section is redundant, it should be removed.

### **Sun 137 Clarify discovery is for initiator**

Editorial

Page 109, F1. The sections in F.1 are involved only in discovery of SCSI peripheral devices by the initiators. The text and titles should be modified to address this.

### **Sun 138 Simplify list**

Editorial

Page 109, section F.1.1, item 7. This item should be divided into two items, like the corresponding items of the list in F.1.2

### **Sun 139 typo**

Editorial

Page 109, section F.2, first line. delete “that”

### **Sun 140 Clarify fabric and device authentication**

Editorial

Page 110, section F.2. The list of items 1-4, is actually two lists, items 1 and 2 addressing the fabric logins and items 3 and 4 addressing the port logins. The text should be separated into two parts. Items 2 and 4 need to be rewritten to clarify the “if-then-else” sense of the sentences. The last part of each sentence (what to do if a configuration change has occurred) needs to be separated out of the respective paragraph and presented as a separate line item or as a separate conclusion.

## Sun 141 Logical Unit Authentication

Technical \*\*\*\*\*

At present, this specifies two device identification page items, port name and node name. This is incorrect. It should be LUN WWN (which may or may not be derived from node name) and optional port name using the association bit.

## Sun 142 Improve informative text

Editorial

Page 111, section G.1. The first sentence, "The required formats for recovery ELSs are described below" s/b "Examples of the formats for recovery ELSs are described below."

## Sun 143 Complete informative text

Editorial

Page 113, Section G.?. Should additional examples be provided for REC and SRR?

## Sun 144 Re-distribute contents of Annex H

Editorial

Annex H should be deleted, and its contents distributed into the body of the document.

Paragraph 1 should be distributed to section 9.4

Paragraph 2 should be distributed to section 9.4

Paragraph 3 should be distributed to section 4.2 or 9.1.

Paragraph 4 should be distributed to a location just before section 4.9.

## Sun 145 Re-distribute contents of Annex I

Editorial

Annex I, with the following modifications, should be moved to section 4.8.

"If a SCSI Target Reset, Logical Unit Reset, or Clear Task Set management function is received by a SCSI Target that has multiple SCSI Initiators logged in with it, then the SCSI Target ~~should~~ **shall**:

- a) create a Unit Attention Condition for all other SCSI Initiators ~~(an FCP\_RSP may have been transmitted but not received by the SCSI Initiator, or the SCSI Initiator may have transmitted a command that has not yet been received by the SCSI Target) (refer to SAM and SAM-2);~~
- b) clear all resources associated with the cleared Exchanges, ~~per SCSI Architectural Model~~ (refer to SAM and SAM-2);
- c) ~~return FCP\_RSP upon completion of (a) and (b). The payload shall be zeroes with the exception of the FCP\_RSP\_LEN\_VALID bit, FCP\_RSP\_LEN (which shall be set equal to 8), and the FCP\_RSP\_INFO (refer to FCP-2).~~ [This is normal behavior already defined for task management, and need not be repeated here].

~~Upon discovery of the Unit Attention Condition set in a), SCSI Initiators should issue ABTS for all commands that are outstanding for the appropriate LUN or LUNs at that SCSI Target as described in 12.5.1. From a SCSI Initiator perspective, this is all commands for which FCP\_RSP has not been received."~~ [This is normal behavior, already covered in other sections.]

**Sun 146 Remove Annex J**

Editorial

These changes to FC-PH-2 and FC-FS should already be in progress and should not need to be covered here.