

PROPOSED RESPONSES TO COMMENTS FOR FCP (X3T9.2/993D)

ROBERT SNIVELY, NOVEMBER 22, 1993

1.0 INTRODUCTION

I have received the following comments on FCP, revision 007 in electronic format during the X3T9.2 review. As I indicated I would in the last meeting, I have analysed those comments for technical and editorial content. For all those editorial comments and for all the technical comments that appear non-controversial, I have described a proposed solution.

For those comments that appear technical and appear to have some controversial content, I have prepared a separate mail that extracts the comments, analysis, and proposed solution from this document to make them easier to consider. Those comments that are discussed in the separate mail are clearly noted in this document.

I assume that all other resolutions that I have suggested are appropriate and will adopt them in the next revision of the FCP document. Notify me as soon as possible if any of these proposed solutions to items that I consider non-controversial are incorrect or controversial, so that I can add them to the separate mail for consideration. Thank you very much for your prompt consideration.

The comments retain their original numbering and are identified by the source of the comments.

\*\*\* indicates that the comment is technical, but probably not controversial. Please review such comments carefully, since they will not be included in the separate document.

2.0 COMMENTS FROM SUN MICROSYSTEMS

Representative: Robert Snively

Date: November 5, 1993

E 001 Cover:

Is "editors", s/b "editor"

Action:

Accepted

E 002 Page i

Is "596-3362", s/b "573-3362"

Action:

Accepted

E 003 Page 2, Section 3.1.5

Is "indicates a displacement", s/b "indicates the displacement"

Action:

Accepted

E 004 Page 2, Sections 3.1.9, 3.1.11, 3.1.14

These sections should be deleted because they are not used by the document.

Action:

Accepted

E 005 Page 2, Sections 3.1.16-3.1.19

These sections should be placed as subsections under 3.1.15 in the same order.

Action:

Accepted

E 006 Page 3

The following terms should be defined on page 3, using definitions paraphrased from SAM:

initiator: an SCSI device containing application clients which originate device service requests and task management functions to be processed by a target SCSI device.

SCSI device: a device that originates or services SCSI commands.

target: an SCSI device containing logical units which receive and execute commands from an initiator.

task: an object within the logical unit representing the work associated with a command or group of linked commands.

Action:

Accepted

E 007 Page 3, Section 3.1.24

The logical unit should be redefined as:

an externally addressable entity within a target that implements an SCSI device model.

Action:

Accepted

E 008 Page 3, Section 3.1.26

The logical unit number should be redefined as:

An encoded 64-bit identifier for a logical unit. [SAM] The logical unit...Address (see section 7.1.1).

Action:

Accepted

E 009 Page 3, Section 3.1.31

The target identifier should be redefined as:

address of up to 64-bits by which a target is identified. [SAM] The target identifier is equivalent to the 24-bit Destination\_ID used by the Exchange Originator. [FC-PH]

Action:

Accepted

E 010 Page 3, Section 3.1.33

The task identifier should be redefined as:

the information uniquely identifying a task. The identifier contains the initiator identifier, the logical unit identifier, and (optionally) the tag. [SAM] The Fully Qualified Exchange Identifier (FQXID) is the task identifier for the FCP.

Action:

Accepted

E 011 Page 3, Section 3.1.34

task management function should be redefined as:

a peer-to-peer confirmed service provided by a task manager that can be invoked by an application client to affect the execution of one or more tasks. The service consists of a request to manage a task in the logical unit. The response.....

Action:

Accepted

E 012 Page 4, Section 4.2

Third paragraph, is "continue to be transferred until all" s/b "continue until all".

Action:

Accepted

E 013 Page 5, Section 4.2

First paragraph, is "or from the SCSI initiator's memory." s/b "or from the SCSI initiator."

Action:

Accepted

E 014 Page 6, Section 5, Table 1

is "Data Delivery Response" s/b "Data Delivery Action".

Action:

Accepted

E 015 Page 6, Section 5

is "maximum burst length by the FCP using the parameters" s/b  
"maximum burst length defined by the parameters".

Action:

Accepted

E 016 Page 6, Section 5

fourth paragraph should be:

If required by the FCP Process Login parameters, each inbound and/or outbound FCP\_DATA IU shall be preceded by an FCP\_XFER\_RDY sequence containing a standard Data Descriptor payload that indicates the exact length of the data delivery. If the FCP Process Login specifies Read XFER\_RDY Disabled and/or Write XFER\_RDY Disabled. The corresponding FCP\_DATA IU's are transmitted without a preceding FCP\_XFER\_RDY IU.

Action:

Accepted

E 017 Page 6, Section 5

is "Some SCSI reset functions are.." s/b "Some SCSI Task Management functions are..."

Action:

Accepted

E 018 Page 6, Section 5, Table 2

Add text to indicate that the Task Management Functions are required to be executed in a new Exchange.

Action:

Accepted

T 019 Page 6, Section 5

\*\*\* Add text to indicate that the CDB is ignored when Task Management functions are executed.

\*\*\* Action:

\*\*\* This technical change is not considered controversial and is accepted.

E 020 Page 7, Section 5, first paragraph

is "..and the Responder X\_ID" s/b "...and optionally the Responder X\_ID".

Action:

Accepted

E 021 Page 7, Section 5, last paragraph

is "...managed by the completion of individual Sequences" s/b "...managed by the transfer of individual IU's"

Action:

Accepted

E 022 Page 7, Section 5.1

last note:

s/b "..optional sequence streaming during Write operations."

Action:

Accepted

E 033 Page 9, Section 5.2

correct minor typos and stylistic errors.

Action:

Accepted

E 034 Page 9, Section 5.2

\*\*\* Add a new sentence, next to the last:

Class 3 service is allowed if the management of the FCP ULP is performed using interlocked IU transmissions.

\*\*\* Action:

\*\*\* Accepted. This comment, while having some technical content, is considered not controversial.

E 035 Page 10, Section 5.3.2 and 5.3.3

Last sentence, is "or striping", s/b "or data striping"

Action:

Accepted

E 036 Page 10, Section 5.3.10

Second and third sentence s/b "The RX\_ID shall be assigned by the Exchange Responder and may have any value, including the unassigned value of hexadecimal 'FFFF'. The RX\_ID allows the Exchange Responder to optionally create a private version of the SAM tag."

Action:

Accepted

E 037 Page 10, Section 5.3.11

Next to last, last, and new last sentence s/b "For the Solicited Data Category, the Relative Offset is the SAM application client buffer offset and the Base Address is the beginning address of the application client's buffer, as described by SAM. For all other ... byte of the IU. The Relative\_Offset is not required if both N\_Ports can unambiguously reassemble the transmitted IU's using other FC-PH information and the FCP\_XFER\_RDY IU."

Action:

Accepted

E 038 Page 11, Section 6

First sentence, second paragraph is "defined as a normative annex" s/b "defined in a normative annex"

Action:

Accepted

E 039 Page 11, Section 6.1

First sentence, first paragraph is "to indicate to" s/b "to identify to"

Fifth paragraph is "for all targets" s/b "for any target"

Last sentence is "from all targets" s/b "from any target"

Action:

Accepted

E 040 Page 12, Table 5

Correct column widths

Action:

Accepted

E 041 Page 13, Section 6.1.1

The word "sequence" should normally be "IU". Correct as appropriate.

The word "will" should normally be "shall". Correct as appropriate.

Action:

Accepted

E 042 Page 13, Section 6.1.1

In both Xfer Rdy Disabled paragraphs, is "FCP\_XFER\_RDY Sequences will not be used.." s/b "FCP\_XFER\_RDY IU's may be not used..."

Action:

Accepted

E 043 Page 13, Section 6.2

is "process images specified", s/b "process image pairs specified"

Action:

Accepted

E 044 Page 15, Section 7.1

The FCP\_CMND IU is actually used for two functions:

- a) Carrying the SCSI command, if and only if no task management function is being requested.
- b) Carrying task management requests, leaving CDB invalid

The text should be clarified to indicate this.

Action:

Accepted

E 044 Page 16, Section 7.1.1

Last sentence should become: "An example of the four-layer hierarchical address structure defined for SCSI RAID systems is given in Annex B."

Action:

Accepted

T 045 Page 17, Section 7.1.2

Should an untagged case be explicitly allowed? SAM defines the untagged case as identical in handling to the Simple\_Q case.

Proposed Action [See separate mail]:

Rejected. Since the SAM untagged case is indistinguishable from the SIMPLE case, only the SIMPLE attribute will be used in FCP. Note that SCSI-2 drivers will be responsible for managing the proper ordering of untagged operations by awaiting the termination of the requested operation, just as they are today. At the same time, tagged operations can flow smoothly in and around the untagged operations, just as in the Queueing Model document.

E 046 Page 17, Section 7.1.2

Clear ACA, third paragraph:

The Clear ACA Task Management function should not be associated with a CDB and is executed in a separate exchange.

Action:

Accepted

T 047 Page 17, Section 7.1.2

Target Reset:

The text should indicate that Target Reset terminates and aborts all outstanding exchanges between the requesting initiator and the receiving target. These events are performed without any FC-PH protocol requirements.

Proposed Action [See separate mail]:

Rejected. The Target Reset does indeed terminate and abort all outstanding exchanges between the requesting initiator and the receiving target. The initiator and target should quit any exchanges that they are aware of and that the other entity has no record of. For those exchanges that are outstanding that must be cleared for proper FC-PH behavior, a sequence using ABTS (last sequence) shall be generated to abort the exchange.

T 048 Page 17, Section 7.1.3

\*\*\* The text should clarify that the CDB is null if a task management function is requested.

\*\*\* Action:

\*\*\* This technical change is not considered controversial and is accepted.

T 049 Page 17, Section 7.1.4

The incorrect length cases should be explained. Only the highest addressed burst of data may be truncated to a length other than the specified burst length. Such truncation will always set the FCP\_RESID residual count and will always set an appropriate FCP\_RESP code. Any bursts of read data beyond the truncated burst will not be transferred. Write operations may not be truncated by the initiator to less than the FCP\_DL.

Proposed Action [See separate mail]:

Accepted. The philosophy is that holes in the data transfer should not be allowed. The additional philosophy is that hosts that are transferring data out from an initiator with a specified data length should have all of that data available to the target if the target wants it all.

E 050 Page 18, Section 7.2

Specifications of transfer size redundant with 7.2.2 or other



sections should be deleted, including first paragraph second sentence and second paragraph first sentence.

Action:

Accepted

E 051 Page 22, Table 14

Correct column widths.

Action:

Accepted

T 052 Page 27 - 29, Annex B

\*\*\* This section should be rewritten to use the RAID addressing example.

\*\*\* Action:

\*\*\* Accepted

### 3.0 QUANTUM

Representative: Jim McGrath

Date: November 5, 1993

E 001 All pages

Quantum is voting NO on forwarding GPP, FCP, and SBP due to the inability to vote YES on SAM. Since these standards must be in compliance with SAM, we do not think it would be consistent to vote to forward these standards without voting to forward SAM.

As a comment, we find the serial protocol documents to have fewer technical issues than SAM, leading us to believe that it might be possible to vote in favor of some or all of them if the SAM issues could be resolved. It is also our understanding that any vote to forward at this point might be ineffective in any event given the pending committee reorganization.

We regret these NO votes, and hope the committee can resolve the outstanding SAM issues as soon as possible.

Action:

None.

### 4.0 NCR

Representative: John Lohmeyer

Date: November 8, 1993

E 001 All Pages

I commend the SAM, SBP, FCP, and GPP document editors on very good draft documents. However, the documents are not yet ready to forward for public review. I feel that it is unfortunate that the committee succumbed to market pressures to issue premature letter ballots on these documents.

I offer the following 6 comments on SAM, with #4 being my reason for voting no. Since SAM is not ready for forwarding yet, I feel compelled to vote no on SBP, FCP, and GPP; it is not appropriate to forward any protocol documents until all the serious issues in SAM are resolved.

Action:

None.

## 5.0 HEWLETT PACKARD

Representative: Jeff Williams

Date: October 15, 1993, modified November 5, 1993

#001 (E) Page 1, Section 2

Add FC-AL to the list of documents referenced.

Action:

Accepted

#002 (E) Page 4, Section 4.1

Add a comment on the ability to connect multiple SCSI devices on an arbitrated loop. I feel that this is an important item for potential implementors of the FCP to understand. SCSI on a peripheral device should be a low cost option which is enabled by the loop topology.

Action:

Accepted

#003 (E) Page 4, Section 4.2

What is an FCP I/O Operation? Is this the same as an I/O Process in the SAM venacular? If so, please eliminate the use of the term in favor of the SAM wording.

Action:

After consulting with the editor of the SAM document, it had been decided that an I/O Process was peculiar to the SAM concepts and did not correctly represent the resources tied to a SCSI command. The term decided upon was an FCP I/O Operation. The intent is to include not only the Command and Task concepts of SAM, but also the Fibre Channel Exchange and all resources associated with it. The editor believes this concept is clear and should not be

changed, but is willing to consider alternative wording.

#004 (E) Page 4, Section 4.2, Para 2

"... and is only used for the transfers associated with that FCP I/O Operation and any commands linked to that command." Drop the "and any commands linked to that command." part. This blurs the definition of the I/O Operation. Is the I/O Operation a single command or a group of linked commands? See the confusion? The definition listed states it is the series of linked commands.

Action:

Accepted

#005 (T) Page 4, Section 4.2, Para 3

You state that the data transfer occurs with a XFER\_RDY IU followed by a single FCP\_DATA IU. Does this have to be the case, or can a target or initiator split the FCP\_DATA IU into two or more IUs for the data described in the XFER\_RDY? For example, I get a XFER\_RDY IU for 8k of data, but I send two 4k FCP\_DATA sequences.

Action:

As defined, the protocol requires and is intended to require that one FCP\_XFER\_RDY IU is followed by one and only one FCP\_DATA IU. The text will be reviewed and modified if necessary to make sure the idea is clearly expressed.

#006 (E) Page 4, Section 4.2

You may want to state that this is a non-linked I/O process with data descriptors enabled. As I read this section, I keep saying, "Sure, but what if the command is linked? Then you won't send this sequence."

Action:

The text in paragraph 4 will be extended and clarified to indicate that a subsequent FCP\_CMND IU may be accepted after the FCP\_RSP IU if all the normal linking conditions are met.

#007 (E) Page 5, Section 4.2, Para 3

List FC-AL as another supported option. You list classes and fabric stuff, so listing loop may be appropriate.

Action:

The text will be modified as appropriate to indicate that FC-AL is an allowable option. See comments 001 and 002.

#008 (E) Page 6, Section 5, Para 3

"SCSI data transfers may be performed by one or more data delivery requests, each one limited to a length no longer than..." This implies that the length of the data delivery request is limited

when you meant that the data delivery length is limited. I would suggest the following wording. "SCSI data transfers may be initiated by one or more Data Delivery Requests, each one limiting the subsequent Data Delivery Responses to a length no longer..."

Action:

Text will be corrected to clarify this point.

#009 (E) Page 6, Section 5, Para 3

Last sentence should state, "... using the parameters in the Disconnect/Reconnect Mode Page."

Action:

The text will be clarified, but it should be made clear that it is the Disconnect-reconnect page of the MODE SELECT command, as specified in 8.2.8, 8.3.3, and 8.3.3.2 of the SCSI-2 document.

#010 (E) Page 6, Section 5, Para 5

"Some SCSI reset functions..." should be "Some SCSI task management functions..." to be consistent with SAM.

Action:

Accepted

#011 (E) Page 7, Section 5, Para 2 (last sentence)

What does this sentence mean? How is the performance of the SCSI command service primitives managed by completion of the individual sequences?

Action:

The wording will be changed to "execution" is "managed by the transfer" of individual IUs.

#012 (E) Page 7, Table 3

Sequences T3, T9 and T10 should have a note that a previous sequence T1, T2, T7 or T8 had to be sent prior to this sequence with the link bit set in the control byte of the CDB. I understand what the intent of these sequences are, but I do not think it is clear to the first time reader.

Action:

A new section between 5.2 and 5.3, entitled "Examples of FCP Sequences" will be created to show how different functions can be supported. Reference to the more detailed examples in Annex A will be included in the section.

#013 (E) Page 7, Table 3, Note 1

IUS should be IUs.

Action:

Accepted

#014 (E) Page 7, Table 3, Note 4

Should this really say T2, T4, T6, T8, and T9 are for streaming? Do I still have to wait for the the ACK on the command prior to the sending of data? If not, I guess I can QUEUE FULL status any command even though the data is received via streaming.

Action:

The examples in the new section, "Examples of FCP Sequences" will demonstrate these possibilities. T2 is only used when initiative is held to perform a subsequent write transfer without FCP\_XFER\_RDY using IU T5 or one or more T6 followed by a T5. Streaming is not allowed for the cases where initiative is transferred, but streaming would be allowed for these cases. This and other examples will be given.

#015 (E) Page 8, Table 4

Are sequences that terminate the exchange (I4 and I6) really doing a transfer of initiative or is this field a "don't care" for last sequences in an exchange? My guess is that we don't care since the exchange has ended.

Action:

The editor sees no guidance in the FC-PH document about this subject. The fields will be changed to X (don't care).

#016 (E) Page 10, Section 5.3.6

Isn't the SEQ\_ID a unique identifier for the sequence so that multiple frames in that sequence can be associated with each other? This states that it is more like the SEQ\_CNT indicating an order to the frames within a sequence.

Action:

Accepted. The section will be updated to conform to FC-PH

#017 (E) Page 11, Section 6.1, Para Last

In the case of a PRLI when one is in effect, what happens to any tasks which are in the device at the time? Are they terminated? Also, I assume that if this affects an initiator/target relationship, that a Unit Attention Condition is created? Please state whether or not this is the case.

Action:

Accepted. Any tasks in the device for the corresponding initiator of the image are removed, as indicated in the paragraph. A Unit Attention Condition may be created if other initiators are effected by the request, but not created if no other initiator is effected. The text will clarify this.

#018 (T) Page 13, Section 6.1, Para 3

\*\*\* For command and data mixed, you state, "If either the originator or the responder do not allow the Command and Data to be combined, then all FCP I/O Operations between them shall be performed using only one Category in each Sequence." I do not agree with this since it also restricts the use of Data/Status combined sequences. This should state that all command and data transfers shall be performed in separate sequences.

\*\*\* Action:

\*\*\* Accepted, ss item 20.

#019 (E) Page 13, Section 6.1, Para 4

"Any other combinations of these bits is invalid and will be rejected." Should say "shall be rejected."

Action:

Accepted

#020 (E) Page 13, Section 6.1, Para 6

See #018. For Read/Status, we should not limit all uses of multiple information categories per sequence.

Action:

Accepted

#021 (E) Page 13, Section 6.1, Para 8 and 10

Replace "wills" with "shall".

Action:

Accepted

#022 (T) Page 15, Section 7.1, Para 2

\*\*\* This paragraph should be replaced with something like,  
"If XFER\_RDY IUs are disabled and Command/Data Mixed is allowed, the Sequence carrying the FCP\_CMND IU shall contain FCP\_DATA IU frames following the FCP\_CMND IU. If XFER\_RDY IUs are disabled but Command Data Mixed is not allowed, sequence initiative may be transferred."

\*\*\* Action:

\*\*\* The requested change is specific, but incomplete, since it does not specify that the behavior is during write operations only. The paragraph will be clarified.

#023 (I) Page 17, Section 7.1.2, Para 5

Terminate Task presents the same problem that Abort Task presented in that the OXID is a duplicate of one that exists. This means that

the interface chip cannot detect duplicate XID creation. Is this an issue that we want to address.

Proposed Action: [See Separate Mail].

Yes, this action clearly needs to be addressed.

Abort Task is managed by the ABTX action or by performing an ABTS with the last sequence bit set.

Terminate Task is managed by the Exchange Originator (Initiator), which performs a Request Sequence Initiative link service request. After the Initiator has gained sequence initiative, it transmits the Terminate Task code in a standard command/task management IU. The task is then terminated according to the SAM rules, requiring at least the presentation of status and sense information describing the completion state of the task.

#024 (I) Page 17, Section 7.1.2, Para 6

Clear ACA also presents a problem. If I want to just clear the ACA and do nothing else, I can't since the command gets executed. I think that this type of operation should state that the CDB payload is considered invalid in some manner. Perhaps this should be done for all task management functions.

Proposed Action: [See Separate Mail].

Accepted. The CDB shall be null for Task Management functions.

#025 (E) Page 17, Section 7.1.2, Para 14

Abort Task is no longer defined. This paragraph should be deleted.

Proposed Action: [See Separate Mail]

The paragraph should be modified to reflect item #023 and included in an explanation of the use of the FCP\_CMND to carry either SAM SCSI service requests or SAM SCSI Task Management functions. It should be treated as a special task management function.

Note that the other task management functions may also require the execution of ABTX or ABTS with last sequence bit set to clear FC-2 activities being carried on for the aborted tasks.

#026 (E) Page 17, Section 7.1.4, Para 1

The FCP\_DL is actually a count of the number of bytes to be transferred for the SCSI CDB, not the I/O Operation (which I don't think we define).

Action:

Accepted

#027 (E) Page 18, Section 7.2, Para 3

The last sentence has some gramatical problems. "...is expected to have be ready to transmit...".

Action:

Accepted

#028 (E) Page 18, Sections 7.2.1 and 7.2.2

Last sentences need periods at the end

Action:

Accepted

#029 (E) Page 17, Section 7.1.4

Add some wording that indicates that insufficient FCP\_DL will cause a command to be stopped in progress. That is, if a write command for 8k were requested, but the FCP\_DL was set to 7k, the command would not complete. Also, what is the expected behavior? Do I write 7k and terminate or just terminate? What about the converse where I have a read for 8k and a FCP\_DL of 7k?

Proposed Action: [See separate mail]

These are all standard incorrect length cases which should be explained in 7.1.4. The following behavior is expected.

A command will begin execution regardless of the FCP\_DL unless the FCP\_DL is zero for a command that requires data transfer.

An incorrect length indication will be presented as described in section 7.4.2 in each of these two cases. If the FCP\_RESID\_UNDER bit is set, the complete data will be transmitted, but the count will be larger than the number of data bytes transmitted. If the FCP\_RESID\_OVER bit is set, the count will be smaller than the number of data bytes that would have been transmitted, but the number of bytes specified by FCP\_DL will be transmitted. The effect on the media depends on the device model being addressed.

#030 (E) Page 21, Section 7.4.3, Para 1

Remove the sentence, "The value is specified in SCSI-2". It has no meaning.

Action:

Accepted

#031 (E) Page 21, Section 7.4.3

Add a sentence that states that the length of this field must be padded to a four byte length due to the restrictions of FC-PH.

Action:

There is no requirement to padd this data to a four byte length. The framing performed by the FC-2 automatically performs this padding for any length.



#032 (T) Page 22, Section 7.4.5

What is the purpose of the FCP\_RSP fields? I see them as additional sense information. Are they considered "soft" error information that presents no reason to retry the command or think that it failed, but gives information useful to the ULP. I see them as a duplication of something already available, namely the sense data.

Proposed Action: [See separate mail]

The FCP\_RSP fields are designed to provide error or status information about the behavior of the ULP, as distinguished from the SCSI task or the FC-2 link errors. The text will be modified to clarify this.

#033 (E) Page 22, Section 7.4.6

\*\*\* Remove references to SCSI-2. I think that we can just say that there may be sense data at any time and that it is in the format shown in the Request Sense command. There is no reason why I can't give sense data on every command.

\*\*\* Action:

\*\*\* Upon re-reading the text, I think it correctly expresses the above concept.

#034a (E) Annex B

Not to throw a monkey wrench into this, but what about a RAID controller attached to a RAID controller, attached to a RAID controller.....

I think this a specific example of something that you would like to do, which is valid, but may not be appropriate for an annex. Find out what the RAID guys have to say and update accordingly.

Action:

Accepted. See Sun comment 052.

#034b (E) Page 4, Section 4.2

The 3rd paragraph under Device management seems awkward. We suggest:

When the device server for a Data-In or Data-Out task has completed the interpretation of the command and is prepared to transmit or receive some or all of the Data associated with the task, it:

1. optionally sends a Data Descriptor IU containing the FCP\_XFER\_RDY payload to the SCSI Initiator to indicate which portion of the data is to follow.
2. sends or receives the solicited data sequence containing the FCP\_DATA payload

Steps 1 and 2 are repeated until all data described by the SCSI command transferred.

Action:

Accepted. The text will be clarified in a manner similar to that requested.

#035 (I) Page 5, Section 4.2

\*\*\* The 2nd to last paragraph has problems:

The design basis for FCP isn't relevant or helpful. The first sentence which describes the wonderful match between FCP and class 2 should be struck. We also believe that class 3 DOES support the FCP with only one possible exception: support of delivery of commands in the order transmitted. Class 3 doesn't exhibit a lack of "confirmed service" support: the passing of Sequence Initiative may count as confirmation for most, if not all of the delivery services.

The paragraph should either be deleted or re-written to say that FCP may be run: class 2-only, mixed class2/class1, or class 3: as determined in a way beyond the scope of this standard (unless you want to suggest that if command comes in class 3, the entire Exchange shall be run in class 3, etc).

If there are specific cases where class 3 does not provide enough confirmation as required by SAM, they should be listed.

\*\*\* Action:

\*\*\* Accepted

#036 (E) Page 6, Section 5

Seems like 4.2 on is the overview. And all of 5 is more detail than overview. Please consider re-titling the section (possibly merge 5 and 6 under 6's title).

Action:

Accepted. Text will be modified to meet the intent of the suggestion.

#037a (T) Page 16, Section 7.1.2

The "Task Identifier" of SAM is missing, and thus untagged commands can't be issued. This breaks the leverage of existing CAM interfaces unless the FCP SIM does some sort of extra management to synthesize untagged behavior on an otherwise tagged-only FCP. The correct solution seems to be to preserve the ability to pass untagged commands. All that is needed technically is to add the "Task Identifier" bit in the command payload; editorial additions clarifying that an OX\_ID tag isn't really a tag if "Task Identifier = Untagged Task Identifier", may also be necessary.

Proposed Action: [See separate mail]

The command queuing annex of SAM states that untagged commands will be treated as if they had the SIMPLE attribute. Since the Fully Qualified Exchange ID is a valid identifier of a task, it should be adequate to map CAM requests for untagged and SIMPLE commands to the SIMPLE attribute as far as

FCP is concerned. The text will be clarified to indicate that.

#037b (T) Page 16, Section 7.1.2

\*\*\* Which of the Task Management Flags should be allowed to be executed concurrently? Our understanding is that SAM presents a single Task Management function per service call paradigm. We suggest that the "flags" should be changed to "functions."

\*\*\* Action:

\*\*\* Accepted. The text will be clarified to indicate that no concurrent Task Management functions can be requested.

#037c (T) Page 16, Section 7.1.2

Also, since SCSI NOP isn't mandatory, we suggest there be a bit that says whether a command is valid. Since it isn't always possible to send Terminate Task, or Abort Task, at least these two must be sent using an FCP Link\_Application look-alike to ABTX: I.e., the OX\_ID must be used to identify the task. However, it must be done on a separate Exchange in order to be able to send it asynchronously. We recommend putting all these functions in a separate Exchange and remove the bits from "normal" commands (except that Clear ACA will appear in both normal commands with ACA\_Q and as an FC4 FCP extended link application):

Task Management flags -> functions (encoding)

reserved		0
Terminate Task	bit 7	1
Clear ACA	bit 6	2
Target Reset	bit 5	3
Clear Task Set	bit 2	4
Abort Task Set	bit 1	5
reserved		:

In summary, we propose the following re-structuring of the FCP\_CNTL field:

bits	Tid	Task Attributes
7   6   5   4   3   2   1   0	-----	-----
+-----+-----+-----+-----+-----+-----+-----+-----+		
MSByte command serial# / res.	0 Untagged	0 Simple_Q
+-----+-----+-----+-----+-----+-----+-----+-----+	1 Tagged	1 Head_of_Q
LSByte command serial# / res.		2 Ordered_Q
+-----+-----+-----+-----+-----+-----+-----+-----+		3 ACA_Q
reserved	R/W	: reserved
+-----+-----+-----+-----+-----+-----+-----+-----+	-----	
ACA Task Attribute  Tid  R/W	0 reserved	
+-----+-----+-----+-----+-----+-----+-----+-----+	1 Read	ACA
7   6   5   4   3   2   1   0	2 Write	----
	3 reserved	0 = Clear ACA
		1 = ACA unchanged

FCP Link Applications:

- Terminate Task
- Clear ACA
- Target Reset

Clear Task Set  
Abort Task Set

Proposed Action: [See separate mail]

SAM presently defines Clear ACA as a Task Management Function, sent independently of a command. I believe that this is correct and that therefore Clear ACA does not appear in "normal" commands.

Comment 37a separately addresses the Tagged/Untagged question.

Comment 37b separately addresses the mutually exclusive nature of the Task Management functions.

In addition, the text should clarify that the presence of any of the Task Management functions should indicate that the CDB is not valid.

The present agreement about the FCP\_CNTL field was reached by negotiation with the SBP editors. There is no strong reason to change it, considering the other solutions.

#038 (T) Page 17, Section 7.1.2 (FCP\_CNTL: Ordered\_Q)

The second sentence describes a procedure which is insufficient to guarantee SAM requirements that commands to a target identifier be received in the order sent. Merely requesting in-order delivery is not sufficient, since guaranteeing in-order delivery involves much more than setting and recognizing the Fabric Login in-order delivery bit.

An alternative is to introduce a serial number in the command: this number would be used to order all commands as appropriate at the device. Such a scheme may require new SCSI errors codes: We haven't researched it (something to indicate: timed-out-waiting for missing command to arrive). For a random access device like a disk, this serial number can probably be ignored unless Ordered\_Q is supported. For sequential access devices, explicit link-independent ordering of commands may be more useful.

The alternative is waiting for some indication of command confirmation. Class 1 and Class 2 offer ACK as a notification that the command was received with good CRC. Although the ACK, doesn't guarantee that the command will be delivered, it is a reasonably reliable confirmation for practical purposes. Class 3 has no such acknowledgment. The command serial number would improve robustness for class 1 & 2 and enable all classes of services to follow the SAM model.

An unpleasant alternative is to wait for confirmation of command received by the target before sending subsequent commands: this requires some sort of polling and/or notification when a command is acknowledged. I.e., subsequent commands to the same LU and/or device will be blocked until each command is acknowledged.

Proposed Action: [See separate mail]

In the case of a well designed fabric that throttles the receipt of frames using BB credit such that F\_BSY never occurs except as a catastrophic error condition, the in-order delivery request is sufficient to make the indicated guarantee.

Therefore the alternative is not necessary.

The command serial number would increase the complexity for no measurable gain. The serial numbers would have to be managed for a very large number of LUN's for commands sourced from a large number of threads, and in the case of hunt groups, across initiator ports. The use of Task Management functions would render the serial numbers non-consecutive, and therefore not useful in managing order. The presence of errors would further complicate the serial number management.

The "unpleasant alternative" cited above is the mechanism used by every operating system that really cares about the ordering of a particular set of operations. The notification function is routinely employed by CAM and potentially has very low overhead.

The suggestion is not accepted.

#### #039 General

We didn't find anywhere that says under what conditions (if any) a Target may provide the FCP\_RSP IU "early". For example, if queue space is limited a Target may choose to go directly to RSP with "Task Set Full" status.

Are there any other restrictions to doing this besides requiring Sequence Initiative? I.e., can it send FCP\_RSP immediately after an FCP\_XFER\_RDY (for a non-zero number of Data-In bytes)?

Action:

These conditions are specified by SAM and needn't be specified in FCP.

#### 6.0 AMDAHL

Representative: Neil T Wanamaker

Date: Nove 5, 1993

001 E Sections 3.1.10, 3.1.12, 3.1.15, 3.1.28

SAM does not appear to require confirmed service here.  
I agree that it is appropriate to use confirmed service here.

Action:

SAM indicates that these are confirmed services, although the confirmation may be no more than the REQ/ACK sequence on a byte by byte basis. No action is required.

002 E 3.1.15.

This refers to the FC\_PH service interface, rather than FC\_PH services.

Action:

Accepted

003 E 4.2., par. 7:

intermix should replace mixed.

Action:

Accepted. The text will be clarified

004 E 5., par.2:

The number of open exchanges is not defined by FC-PH characteristics, but rather negotiated at login as defined in FC-PH. The number of exchanges is for the pair of N\_Ports, not qualified by FC-4 type.

Action:

Accepted. The text will be clarified

005 E 5., Table 2:

Functional negotiation isn't in SAM. Login/Logout isn't optional. (note typo in header).

Action:

The observation is correct. The text will be modified to separate link level functions from the SAM functions.

Login/Logout need not be performed for those systems using implicit login techniques.

006 E 6.2.1

First sentence should refer to "Logout".

Action:

Accepted

007 E 7.1

second sentence: "..an.. frames..."

Action:

Accepted

008 E 7.2

last sentence of par.3: "..have be ready"

Action:

Accepted

009 E 7.2.1.

\*\*\* What is the initiator to do if the relative offset of the data frame doesn't match the Transfer Ready?

\*\*\* Action:

\*\*\* This is an FCP ULP error and, if detected, should be posted as a new response code in Table 14, "ULP Data Transfer violation".

## 7.0 IBM

Representative: George Penokie

Date: November 7, 1993

001 E Section 4.2 paragraph 2,

change UI to IU.

Action:

Accepted

002 E Section 6.1 paragraph 1:

I believe should not refer to a "command pair." I think this should just say "command."

Action:

The request/response pair for PRLI requires IUs from each of the two devices creating an image pair. The text will be clarified to indicate this behavior.

003 E Section 6.1 paragraph 7

This section refers to a PRLI to an "image pair," while paragraph 1 indicates PRLI goes from an originator N\_Port to a Responder N\_Port--not an image pair. Perhaps what is said is correct but I think it needs clarification.

Action:

Accepted

004 T Section 6.1

This section should probably state that all SCSI commands associated with a process which logs out are aborted. It seems unclear from reading the document whether SCSI commands are able to span process logout/login cycles.

Action:

Accepted. Note that the other task management functions may also require the execution of ABTX or ABTS with last sequence bit set to clear FC-2 activities being carried on for the aborted tasks.

005 7.1.1

This section states that the entity address is vendor unique. If

possible it would be better to refer to the RAID document on addressing, X3T9/93-161. Annex B, although informative, should also be updated to conform to that document.

Action:

Accepted. See Sun Comment 52.