Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900

June 4, 1999 T10/99-211 revision 0

To: T10 Membership
From: Bob Snively, Sun Microsystems
Subject: FCP-2 changes between revision 01 and revision 02

The following changes have been placed in FCP-2 between revision 01 and revision 02. These changes have been approved in several technical meetings since the publication of revision 01. In addition, small editorial errors and technical clarifications that have been called to my attention and discussed by e-mail or in the committee have been corrected.

Technical changes

Section	Ref	E/T	Summary	Approved
4.4	1.10	T	FCP_CONF and linked commands	4/15/99
5.1	6.2	T	Make operation associators obsolete	FC-FS
5.3	1.6	T	Obsolete mixed command/data, data/response	4/15/99
6.2.6	1.12	T	Determination of FCP recovery capabilities	6/10/99
6.2.7	8.4	T	Behavior of PRLI	N/A
7.1	2.1	T	SRR should be moved	e-mail
7.1	2.2	T	Correct reason code	e-mail
7.1	2.3	T	Correct reason code	e-mail
7.1	2.4	T	Verify reason codes	e-mail
7.1	1.4	T	SRR Reason Code incorrect	4/15/99
7.1	1.9	T	SRR to be rewritten	4/15/99
8.1	1.2	T	Long CDB	6/10/99
8.1.1.4	6.1	T	Mandatory and optional task management	e-mail
8.4	1.8	T	Interpretation of FCP_RSP	4/15/99
9.1.1	2.5	T	Clarification of Buffer Empty Ratio	e-mail
9.1.1	1.7	T	Clarify First Burst Length parameter definition	4/15/99
9.1.3	1.1	T	RR_TOV timer definition	4/15/99
9.1.3	1.11	T	Private mode bit	6/10/99
10.6	1.5	T	REC_TOV timer implementation	4/15/99

Technical changes

Section	Ref	E/T	Summary	Approved
11	2.6	T	Clarification that link error recovery works if in-order	e-mail
11.2.4	2.7	T	Clarify case of managing sequence initiative	e-mail
11.2.4	1.13	T	Discarding of Exchange Status	5/4/99
11.2.4	1.14	E	Discarding exchange information	e-mail
11.3	2.8	T	Correct error recovery case	e-mail
11.4.2	2.9	T	Allow reject if unknown FQXID received	e-mail
B.2.1	5.1	T	References for ABTS	5/4/99
B.2.1	6.3	T	Discovery of ABTS capability	N/A
B.3.1	1.3	T	REC reject reason code corrected	4/15/99

Editorial changes

Section	Ref	E/T	Summary	Approve
All	3.4	E	Parameters in small caps	NA
Intro	3.1	E	Correct clause descriptions	NA
All	4.1	E	Reorder annexes	e-mail
3.1.7	4.2	E	Correct definition of CDB length	e-mail
3.3	3.9	E	Add clause for keywords	NA
4.1	4.3	E	Count of classes of service	e-mail
4.1	4.4	E	Loop attached fabric correction	e-mail
4.1	4.5	E	Table 1 incomplete	e-mail
4.2	4.6	E	Description of FCP_CONF required	e-mail
4.2	4.7	E	Definition of classes of operation	e-mail
4.4	4.8	E	Incorrect reference	e-mail
4.4	4.9	E	Clarification of FCP_CONF	e-mail
4.4	4.10	E	Usage of FCP_CONF	e-mail
4.4	4.11	E	FCP_CONF retransmission	e-mail

Editorial changes

Section	Ref	E/T	Summary	Approve
5.1	4.12	E	Glossary terms for OOA and ROA	e-mail
5.4	4.13	E	XFER_RDY on reads is optional	e-mail
5.5	4.14	E	Clarification of classes of service	e-mail
6.2.7.2	4.15	E	PRLI reference	e-mail
7	4.16	E	typo	e-mail
7	3.10	E	Remove editor's note	NA
7.1	4.17	E	SRR sent by initiator only	e-mail
8.1.1.4	3.7, 4.18	E	Correct insertion of table	NA
8.1.1.6	4.19	E	CDB size	e-mail
8.1.2.2	4.20	E	Verify ABTS with unassigned RX_ID	e-mail
8.2	4.21	E	Grammar	e-mail
8.3	4.22	E	Incorrect use of SEQ_ID	e-mail
8.3	3.8	E	Correct clause 8.3	5/4/99
8.4	3.2	E	Warning of legacy FCP_RSP definitions	6/10/99
9.1.1	4.23	E	Missing information in column	e-mail
9.1.3	4.24	E	Clarification of attachment to point-to-point link	e-mail
10	3.3	E	Redundant timer definitions removed	4/15/99
11	3.6	E	References for error examples	NA
11	3.5	E	Cross reference modifications in clause 11	NA
11.1.1	4.25	E	Clarify case for missing responses	e-mail
A	4.26	E	Arrows needed on ACKs	e-mail
C	4.27	E	Underlining incorrect	e-mail
C.5	4.28	E	Exchange completion clarified	e-mail
C.6	4.29	E	Correct error recovery case	e-mail

Proposed changes that are not installed

Section	Ref	E/T	Summary	Source
4	7.1	T	Additional discovery proposals	e-mail
10	7.2	T	Default E_D_TOV	
11.2.7	7.3	E	Target never sends REC	e-mail
B	7.4	E	Change location of SRR	e-mail

1.0 Technical Changes

1.1 RR_TOV timer

The RR_TOV timer is now installed in bytes 6 and 7 of page 19h in section 9.1.3 of the FCP-2 document. The format and location of the control bits has been changed as a result of analysis by the working group. The following text and modified tables are included in 9.1.3.

Fibre Channel Port Control page (19h)

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Resvd	Page Code (19h)					
1	Page Length (06h)							
2	Reserved							
3	Resvd	PLPB	DDIS	DLM	DSA	ALWI	DTIPE	DTOLI
4	Reserved							
5	Reserved							
6						RR_TOV units		
7	Resource Recovery Time Out Value (RR_TOV)							

The RR_TOV (see TBD) is defined by bytes 6 and 7 in the following manner.

The RR_TOV units field indicates the units in which the RR_TOV is calculated, according to table x.

Table x: Values for RR_TOV units

Byte 6			Units of measure for RR_TOV
bit 2	bit 1	bit 0	
0	0	0	No timer is specified
0	0	1	0.001 seconds
0	1	1	0.1 seconds
1	0	1	10 seconds
All other values			Reserved

The RR_TOV field indicates the number of time units specified by the RR_TOV units field that shall be used by the timer that performs the RR_TOV timeout functions. If no timer is specified, the RR_TOV value in byte 7 shall be ignored by the device server. Those functions are specified by FC-PLDA and by section TBD of this standard.

The supplementary text for this timer has been approved by the actions in 3.1 of this letter for inclusion in revision 02 FCP-2.

1.2 Long CDB

The DL field had been placed after the FCP_CDB field, making the FCP_CDB and “additional FCP_CDB” fields contiguous. The position of the FCP_DL field in a long CDB (clause 1.5) was discussed many times. A vote was finally taken and the group voted 12 to 0 to put the FCP_DL field at the end of the FCP_CMND. The current wording in FCP-2 will be retained. (May 4)

1.3 REC reject reason code corrected

Section B.3.1 was modified to correctly reference the reason codes for rejection of REC.

Section B.3. 1 was corrected to clarify the meaning of the Data Transfer Count field in the REC Accept payload.

Neil Wanamaker has provided a correction of the reason code in an e-mail dated April 28, 1999. The code is corrected to be 09/2A, indicating the reject reason code of “unable to perform command request” and “unable to supply requested data”.

1.4 SRR Reason Code incorrect

Section 7.1 was modified to correct the SRR LS_RJT reason code, 00052A00h and found that it is mislabeled.

Neil Wanamaker has provided a correction of the reason code in an e-mail dated April 28, 1999. The code is corrected to be 09/2A, indicating the reject reason code of “unable to perform command request” and “unable to supply requested data”. Additional reason codes are included by reference.

This correction will be put in revision 02 of FCP-2.

1.5 REC_TOV timer implementation

The timer definition is softened in clause 10.6 to require that at least REC_TOV shall pass before making the first REC for an exchange. That allows implementation with a single timer plus a single control bit for each exchange, but forces no particular implementation. The change was made in revision 02 of FCP-2.

The committee determined that a further assumption has to be made to make this operate correctly. RR_TOV shall be greater than or equal to 3 REC_TOV if REC_TOV recovery is used. Apparently, consideration of ADISC address discovery time must also be included. The following note is added to clause 10.5

> If REC_TOV associated error recovery is allowed, RR_TOV must be 3 times REC_TOV and always appropriate to ADISC address discovery time.

1.6 Obsolete mixed command/data and data/response

The FCP-2 document is simplified by removing and making obsolete the use of mixed command/data and mixed data/response. That:

a) eliminates some text in section 8.1,

b) removes the necessity to clarify sections 8.3 and 8.4,

c) eliminates the PRLI parameters of command/data mixed allowed and data/response mixed allowed in 6.2.6 Table 9, 6.2.6.11, 6.2.6.12, and 6.2.7 Table 10, and

d) removes IUs T8, T9, T10, T11, I6, and I7 in section 5.4.

It would saves the effort of resolving and documenting the error recovery procedures for these cases.

1.7 Clarify First Burst Length parameter definition

In section 9.1.1, the definition of the First Burst Size value will be clarified. The present text indicates that First Burst Size is related to tenancy. The text will be modified to indicate that the First Burst Size has nothing to do with tenancy. The words “tenancy” and “loop” were searched and their relevance reviewed throughout the document. Minor wording changes were installed in clauses 9.1.1, 9.1.1.2, and 9.1.1.3.

The text was modified to indicate the conditions under which the first burst is limited to the value specified by the First Burst Size field. Those conditions are:

1) The burst size applies only for operations with outbound data transfers, and

2) The burst size applies only if Write XFER_RDY is disabled.

Under all other conditions, the value of the First Burst Size field will be ignored.

In addition to these changes to section 9.1.1.10, headers were provided for each bit.

1.8 Interpretation of FCP_RSP

At present, it is not always clear that an FCP_RSP_LEN_VALID or an FCP_SNS_LEN_VALID bit of 0 means that the corresponding length field is to be ignored, regardless of its value. This additional clarification will be added to clause 8.4 revision 02 of FCP-2.

The text was additionally corrected to clarify that “ignored” is interpreted as having a length value of zero.

1.9 SRR to be rewritten

At present, there is an attempt to define SRR as a generic recovery mechanism. The attempt quickly slips into direct SCIS FCP definitions when the individual fields are defined. I have already changed it to an FC-4 link data frame. Section 7.1 was rewritten to reflect the SCSI-specific use of SRR as an initiator-only function.

1.10 FCP_CONF and linked commands

Linked commands were not allowed to request FCP_CONF. No change will be made.

The committee requested that the document allow FCP_CONF on the terminating CDB of a set of linked commands, whether the termination is caused by reaching the end of the set of linked commands or by encountering a check condition that ends the linking process. In addition, the text should explicitly prohibit the use of FCP_CONF on other linked command cases. The text of revision 02 of FCP-2 was modified to read:

If command linking is being performed, the target shall not request confirmed completion. The target may request confirmed completion

a) when providing FCP_RSP for the last command of the set of linked commands, or

b) when providing FCP_RSP for a command that terminates linking because of an error or CHECK CONDTION status.

1.11 Private mode bit

A “private mode only” bit was approved for inclusion in FCP-2. The bit is documented by 99-319v1. The changes are installed in table 26 of 9.1.3, in 9.1.3.6, and in 9.1.3.8.

1.12 Determination of FCP Recovery capabilities

There have been several discussions about how the more sophisticated functions of FCP-2, including execution of REC, execution of SRR, support of FCP_CONF, and support of CRN are discovered and enabled. After considerable discussion, the committee during the June 4, 1999 meeting has recommended the following discovery controls:

The committee has elected to put FCP_CONF in the PRLI. That is documented in FCP-2, rev 01, clauses 4.4 and 6.2.6.7

The committee has elected to put CRN in the MODE SELECT page for Fibre Channel LU control. That is documented in FCP-2, rev 01, clauses 4.3 and 9.1.2.

The committee has chosen to create an additional bit that indicates support for SRR in the PRLI. The bit, when present, also indicates that REC is supported. This is to be documented for the first time in FCP-2, rev 02. The case where data is unavailable to perform the retry requested by the SRR ELS is already documented by FCP-2, rev 01. Note that some logical units in a target may never be capable of performing such a retry.

REC is assumed to be supported whether the SRR PRLI control bit is on or off, although the target may choose to reject the REC if it is not supported. Note that REC may even shorten recovery actions for disk drives. This is to be documented for the first time in FCP-2, rev02.

The use of REC and SRR is summarized in a new paragraph in section 4.

The discovery mechanisms for FCP-2 capabilities is summarized in a table in section 4.

The definition of the PRLI bit that controls SRR is placed at the proper location in clause 6.2.6. The assigned bit (word 3, bit 8) is named “Retry” since it indicates the presence of the transmission retry capability invoked by SRR. Both the initiator and target must have the Retry capability indicated for SRR to be used. If both have the Retry capability, SRR shall be used according to the rules in FCP-2.

The “Retry” bit is also added in 6.2.7, table 10, since the target must agree to support the functionality. Note that the other bits (Confirmed Completion Allowed and Data Overlay Allowed) are capabilities announced by the initiator to the target and are optionally exploited by the target if available. No ACC is required.

1.13 Discarding of exchange status

Dave Peterson made a comment that FCP-2 needs to include text stating that in a class 3 non queuing environment exchange status can be discarded after RR_TOV as well as after receiving another command from the same initiator. After a discussion with Dave, Bob and Stewart Wyatt, Bob agreed to make the change. (May 4, 1999)

See also 1.14

1.14 Discarding exchange information

11.2.4 - Isn't the target allowed to discard exchange information after RR_TOV? (Question from Matt Wakeley)

2.0 Technical corrections from HP

The following corrections were received by e-mail from Matt Wakeley. They were discussed by e-mail and the proposed corrections were accepted without dissension.

2.1 SRR should be moved

7.1 - SRR is a new “FC-FS” feature - shouldn't it be moved to Annex B where REC and the new ABTS features are documented?

Rejected, SRR is an FC-4 ELS. See 1.9.

2.2 Correct reason code

7.1, second paragraph - the reason code (05) does not match “unable to perform the command request” - this is code 08. 05 is defined in FC-PH as “logical busy”. Also, what is the reason code explanation code for the indicated error? Accepted. See minutes of meeting and info from Wanamaker, section 1.4.

2.3 Correct reason code

7.1, third paragraph - the reason code (0005A200) should probably be 00082A00 (08, not 05 and 2A not A2), see previous comment. Accepted. See minutes of meeting and info from Wanamaker, section 1.4.

2.4 Verify reason codes

7.1, table 15 - the reason codes are probably wrong, and there should be at least two. References to FC-PH are provided for any other reason codes that may be required. For the specified reason code, see minutes of the meeting and info from Wanamaker, section 1.4.

2.5 Clarification of Buffer Empty Ratio

9.1.1, page 44 near bottom of page. The buffer empty ratio field during a write operation is supposed to define how empty the buffer should be prior to requesting an interconnect tenancy. How is the device server (target) going to tell the initiator when the initiator can or cannot arbitrate for the interconnect? The only way it could do this is via not sending the FCP_XFER_RDY. Your interpretation is correct. The text is modified in 9.1.1 to clarify this.

2.6 Clarification that link error recovery works if in-order

11 - It needs to be specified, in big bold letters, that the link error recovery procedure specified here ONLY WORKS ON AN IN-ORDER TOPOLOGY.

Accepted

2.7 Clarify case of managing sequence initiative

11.2.4 - the test is for “E_STAT indicates that the Exchange is complete”. However, if the target requested a FCP_CONF, the exchange will not be complete. The test should be “E_STAT indicates that the exchange is complete or the initiator holds SI”. However, this additional test of “initiator holds SI” can cause the initiator to get confused. It won't know the difference between the FCP_DATA (write) being lost, or a FCP_RSP with FCP_CONF_REQ being lost. In both cases, the initiator will hold SI. The initiator will also have to check the data received count in the REC. This also applies to figure C.5.

Section 11.2.4 was rewritten to correct this, but should be reviewed very carefully. The figure in C.5 (now D.5) was relabeled to indicate that it was a case where FCP_CONF was not included as a possible action.

2.8 Correct error recovery case

11.3 - The error recovery specified may cause the target to abort the wrong exchange. Consider the following example: The target sends FCP_RSP (class 2) with the end of exchange bit set. The initiator responds with an ACK to the FCP_RSP with the end of exchange bit set. As far as the initiator is concerned, the exchange is terminated. Now, if the ACK gets lost on its way back to the target, the target still has the exchange alive. The initiator is free to reuse the OX_ID for that exchange. If the initiator starts a new exchange

with the same OX_ID at the same time the target decides to send an ABTS for the ACK it never got back, the target is now aborting the new (wrong) exchange instead of the old one where it lost the last ack. (remember, the target is not required to assign an RX_ID, so the initiator may not be able to figure out that the ABTS is for a long since dead exchange.)

The solution is that for class 2, the target must always request FCP_CONF, or must never send an ABTS for a missing ACK to an FCP_RSP.

Instead, the following wording was selected. Please review this solution.

When OX_ID values are reused quickly and RX_ID values are not used, it is possible for a missing ACK for FCP_RSP to allow the target to attempt to abort a more recent exchange using the same OX_ID. To prevent that, a target using class-2 behavior shall either request FCP_CONF or shall use RX_ID to distinguish outstanding exchanges.

2.9 Allow reject if unknown FQXID received

11.4.2 - the target should be allowed to send a BA_RJT if an ABTS with an unknown FQXID is received.

Accepted. Case B in the clause already attempts to cover this, but becomes incomplete by focusing on the period before an RX_ID is exchanged. Case B was modified to include any FQXID.

3.0 Editorial corrections

3.1 Correct clause descriptions

In the Introduction, the clause descriptions were updated to reflect the present structure of the document.

3.2 Warning of legacy FCP_RSP definitions

While no change is required in the technical text of revision 02 of FCP-2, the committee recommended that a note indicating the existence of older non-conforming 12-byte fields be included in clause 8.4

3.3 Redundant timer definitions removed

Those timers defined for functions not described in FCP-2, including AL_TIME and R_T_TOV will be removed from table 28 of clause 10 in revision 02 of FCP-2. Sections 10.1 and 10.2 will similarly be removed.

3.4 Parameters in small caps

All sections:

Parameters should be identified by small caps instead of full-size caps. This correction is partially installed in revision 02 of FCP-2. The final set of changes will be in revision 03 of FCP-2.

3.5 Cross reference modifications in clause 11

Section 11.2.8 points to an FC-TAPE definition of ULP_TOV timeout with an incomplete exchange. The reference is deleted, since the timers are not placed in FCP-2 yet. Deleted text:

a) ULP_TOV Timeout and Exchange Not Complete (see 5.13 for conditions indicating Exchange completion)

The remaining cross references can now be updated, since I have the new versions of this chapter. In addition, I have pulled in the timer definitions not covered by other standards.

These corrections will now be completed in revision 02 of FCP-2.

3.6 References for error examples

To assist in understanding the error recovery procedures, cross references between the error recovery examples in Annex C and the error recovery descriptions in clause 11 are provided. So far, only the general cross-reference has been completed. I will work on these as time permits, but it will not be complete until FCP-2 revision 03.

3.7 Correct insertion of table

Table 18 of section 8.1.1.4 was updated incorrectly in revision 01 of FCP-2. The table should contain the required task management flags. This correction is installed in revision 02 of FCP-2.

3.8 Correct clause 8.3

In clause 8.3, a reference to a “technical report” was changed to be a reference to a “standard”.

3.9 Add clause for keywords

Clause 3.3 is expanded to include the definition of keywords required to understand standards.

3.10 Remove editor's note

Clause 7 has an editor's note that merely requested review. It is now removed.

4.0 Editorial corrections from HP

The following corrections were received by e-mail from Matt Wakeley. They were discussed by e-mail and the proposed corrections were accepted without dissension.

4.1 Reorder annexes

page xiii and the Annexes. I think it is customary to put all the "normative" annexes before the "informative" ones. Accepted

4.2 Correct definition of CDB length

3.1.7 - the definition of "command descriptor block" still indicates it is up to 16 bytes in length, which conflicts with the "long CDB" and 8.1. Accepted

4.3 Count of classes of service

4.1 - the first paragraph describes "three" classes of service. The reference to how many classes of service there are defined in FC should be removed. Accepted.

4.4 Loop attached fabric correction

4.1 - the third paragraph describing FC-AL should also indicate that communication can be performed between an NL_Port and an FL_Port for communication to a loop attached fabric and other ports attached to the fabric. Accepted in principle, but the wording must be reviewed.

4.5 Table 1 incomplete

4.1, table 1 - FCP_CONF is missing from the table. Accepted

4.6 Description of FCP_CONF required

4.2, page 8, first complete paragraph, describing FCP_RSP - There probably needs to be a description of FCP_CONF here. A reference was provided to the description of confirmed completion in the fifth paragraph of 4.2.

4.7 Definition of classes of operation

4.2, page 8, 5th paragraph describing classes of operation - References to class numbers should be changed to “acknowledged” and “unacknowledged” classes of service.

Accepted.

This paragraph is obsolete. I replaced it with the following text:

FCP-2 takes full advantage of the multiplexing and shared bandwidth capabilities provided by various Fibre Channel classes of service. The protocol is designed to operate with any class of service and to provide options for reliable error detection and error recovery independent of the class of service.

4.8 Incorrect reference

4.4, first paragraph - The reference to FC-Tape should be changed to FCP-2, since the retry mechanisms have been moved to this document. Accepted.

4.9 Clarification of FCP_CONF

4.4, second paragraph - Change the sentence “If the FCP_CONF is not returned the SCSI target can retransmit...” to “... the SCSI target may be requested by the initiator to retransmit...” Accepted.

4.10 Usage of FCP_CONF

4.4, page 10, first paragraph - “...FCP_CONF confirms that the FCP_RSP of GOOD status has been received...” delete “of GOOD status”. FCP_CONF is not relegated to only being used for “good” status - it can also be used to confirm the delivery of “error” status. Accepted

4.11 FCP_CONF retransmission

4.4, page 10, first paragraph - “If the FCP_CONF is not returned, the SCSI target can retransmit...” change to “... the SCSI target may be requested by the initiator to retransmit...” Accepted

4.12 Glossary terms for OOA and ROA

5.1, table 4 - the terms “OOA” and “ROA” are not defined in the document. They have been deleted. See 6.2.

4.13 XFER_RDY on reads is optional

5.4, table 7 - I2 is defined as mandatory. However, most FCP devices don't implement XFER_RDY on reads, so this should be marked “optional”. Accepted.

4.14 Clarification of classes of service

5.5, first paragraph - change “Class 1 and Class 2 service provides...” to “Acknowledged classes of service provide...”. Accepted.

4.15 PRLI reference

6.2.7.2 - The PRLI accept response code is now defined in FC-PH-2, not Annex A. Accepted.

4.16 typo

7, second paragraph - typo “tT”. Accepted.

4.17 SRR sent by initiator only

7.1, first paragraph - in FCP, the SRR is only sent by the initiator (exchange originator) to the target (exchange responder), so the references to the responder sending it should be removed. Accepted.

4.18 Correct table 18

8.1.1.4, Table 18 - table 18 is an exact copy of table 17. It needs to be changed to include the correct information. Accepted.

4.19 CDB size

8.1.1.6 - the FCP_CDB is defined as a maximum of 16 bytes (unless a long cdb is used). It seems to me that the CDB is always exactly 16 bytes (unless a long cdb is used) so that the FCP_DL field is always in the same place. No change is required. FCP_CDB is always a specified size, but the CDB is not.

4.20 Verify ABTS with unassigned RX_ID

8.1.2.2 - “A target shall always accept an ABTS using the unassigned RX_ID value of “FFFF” hexadecimal and establish a recovery qualifier with a specified RX_ID.” Shouldn't this say “OX_ID”?

I have reviewed this in context and believe that it is correct as written. The intent appears to be to allow the recovery abort of an exchange where the FCP_CMND has been successfully transferred, but no returned frame has yet provided the RX_ID to the initiator. As a result, the ABTS will use the original OX_ID and the non-addressed RX_ID value. The ABTS ACC will contain the actual RX_ID which was to be used by the exchange.

No change is required.

4.21 Grammar

8.2, second paragraph - typo “after the first of one of a” - delete the first “of”. Accepted.

4.22 Incorrect use of SEQ_ID

8.3, bottom of page 36 - “...the SEQ_ID and the RLTV_OFF are used to ensure that the SCSI data is reassembled in the proper order.” How is “SEQ_ID” used to ensure “orderness”? If anything, it should be SEQ_CNT, and that only works if SEQ_CNT is continuously increasing. Accepted. RLTV_OFF is the only value required, once proper exchange identification has been performed.

4.23 Missing information in column

9.1.1, table 24 - the byte numbers 8-15 in the left hand column of the table are missing. Accepted.

4.24 Clarification of attachment to point-to-point link

9.1.3, change all occurrences of “Targets attached to an N_Port or to an F_Port shall ignore this bit.” to “Targets not attached to an FC-AL Loop shall...” to be consistent with the requirement “...a target attached by an FC-AL loop shall...” Accepted

4.25 Clarify case for missing responses

11.1.1, item 2 should be changed to: “... and no FCP_RSP IU or FCP_XFER_RDY IU has been received for the FCP_DATA IU.” Accepted.

4.26 Arrows needed on ACKs

Annex A - figures A.1 - A.4 need to add arrows on the ACKs. Accepted.

4.27 Underlining incorrect

Annex C - all text that looks like it was intended to be underlined, is instead “under arrowed”. Accepted. This appears to be a frame-to-pdf bug. Substituted bold type for underlined type.

4.28 Exchange completion clarified

Figure C.5 - see my comments to 11.2.4 (in this document, 2.7). The test only for the exchange being complete is not sufficient. Needs to be reviewed and clarified as above.

See 2.7.

4.29 Correct error recovery case

Figure C.6 - see my comments to 11.3 (in this document, 2.8).
See 2.8.

5.0 Items reviewed by editor

5.1 References for ABTS

FC-TAPE review comment 31 from Charles Binford provided references for the ABTS ELS, which would require changes to FC-TAPE. The function was removed from FC-TAPE and moved to FCP-2 Annex B, section B.2.1. However, only the modified portion was documented there and the remainder was included by reference, so no changes should be required. The editor's note in this section is now removed. The editor's review of this section required a small change to section B.2.1.

At Charles Binford's suggestion Bob will check and repair references to ABT-LS for recovery. The new kind of ABT-LS is the recovery type. Bob will place a note near ABTS-LS recovery abort to clarify the changes made to ABT-LS.(May 4)

After review of B.2.1, 11.4, and 8.1.2.1, no changes were made.

5.2 Rules for ELS generation before Login

Most FCP devices compliant with FC-PLDA limit the ELS codes that may be used before a Login has been successfully completed. At present, FCP is silent on this and FC-TAPE has expressed rules similar to FC-PLDA. Robert Kembel's comment #29 on FC-TAPE, which requires clarification on this issue, has not been resolved.

The committee has requested Bob Snively and Bob Kembel to prepare a list of ELSs that do not require implicit or explicit login. This list, when approved, will be included in an annex of revision 02 of FCP-2 until it is transferred to FC-FS.

Bob Kembel indicates:

I was wondering which Extended Link Services that you feel should require PLOGI before they are accepted. The only one that I would accept without argument is PRLI. Are there others?

This work item is not yet complete and should be included in FCP-2, revision 03.

5.3 Consideration of additional ELSs

Ralph Weber asks if ADISC, PDISC, and FDISC should be added to the extended link services documented in FCP-2. I believe the present structure, which tells how FCP-2 relates to the ELSs documented in other locations is probably still a good documentation method.

Dave Peterson accepted an action item to document the discovery protocol for inclusion in a standard to be determined, probably either FC-FS or FC-FA. The June, 1999 FC-TAPE meeting created the action item of placing the discovery description in an FCP-2 annex for now.

There is still an editor's note requesting further input on these and other questions.

1. Discuss ADISC/PDISC/FDISC/FAN? (Add PLOGI)
2. Comment: HBA drivers do not issue INQUIRY or REPORT LUNS done by Class driver (maybe).
3. Comment: Track each SCSI device (i.e. LUN) using the WWNN or WWPN returned in the INQUIRY or PLOGI.
4. Comment: Add text regarding REQUEST SENSE.
5. RNC usage?

6.0 Technical changes to be approved

The following technical changes, although approved in various meetings and by e-mail, are relatively complex and deserve special review and approval.

6.1 Mandatory and optional task management

8.1.1.4, bottom of page 33, Abort Task Set - this bit is indicated as “mandatory by FCP”. What other bits are “mandatory” and which are “optional”?

The intent of the question was accepted. The following table indicates my best interpretation of the requirements.

Table 1 - TASK MANAGEMENT Flags

bit	TASK MANAGEMENT FUNCTION	MANDATORY/ OPTIONAL
7	obsolete	
6	CLEAR ACA	Mandatory if NormACA bit in INQUIRY set to one. Prohibited if NormACA bit in INQUIRY set to zero.
5	TARGET RESET	Mandatory
4	LOGICAL UNIT RESET	Mandatory
3	reserved	
2	CLEAR TASK SET	Mandatory
1	ABORT TASK SET	Mandatory
0	reserved	

6.2 Make OOA and X_ID invalidation obsolete

The FC-FS ad-hoc has agreed to make OOA obsolete. While I still have to provide specific proposed text for that to the FC-FS group, I have already removed the references to OOA in clause 5.1. The glossary entry 3.1.25 is also deleted.

6.3 Discovery of ABTS capability

Robert Reynolds asks about the new format specified for the ABTS command, where the Parameter field determines whether to do a sequence or exchange level abort. How does an Initiator determine target compliance for this new ABTS format?

This is now clarified in B.2.1. Please review this.

7.0 Technical proposals not included

7.1 Additional discovery proposals

Robert Reynolds presented the following discussion about determination of FCP-2 node compliance.

“There is a mechanism to determine support for individual features of FCP-2 but not for overall compliance. There is nothing specified in the PRLI parameters and the FC-PH revision levels are the same as for FLA.

For most of the features a node can determine compliance by checking either the PRLI parameters or use the new Mode pages. However, there are two new features that I don't see how to determine if they are supported by a node.

With the latest proposal, an FCP_CMND with additional CDB data present would look identical to an FCP_CMND in the FCP-1 format. How does an Initiator determine whether it can send the new command format? Is there going to be a new PRLI bit added to specify this support? Without some way to determine this a FCP-1 compliant target would receive the new command format and process it as if it was the old format.

Robert Elliot explained that the new version descriptors in the INQUIRY data (see SPC-2 revision 9 or later) can be used to indicate whether the device was designed to comply with FCP or FCP-2 (and possibly identify specific revisions of each).

While the discussion continued, the question basically came down to two issues:

a) FCP link capability discovery

b) Long CDB capability discovery.

The first is now addressed by the discovery negotiation described in 1.12.

The second is not an FCP problem, but rather a SCSI problem. It will most likely be addressed by defining a device type that exploits the new capabilities. An example of such a new device type would be an object oriented disk, which is not presently defined by any standard.

No change was made to FCP-2 to address this question.

7.2 Default E_D_TOV

Dave Baldwin indicates that FCP-2 Table 28 says E_D_TOV footnote 3 on page 49 places the fabric/pt-pt default as 10 seconds.

He did not find a justification for this in any document, and would prefer to see 2 seconds.

7.3 Target never sends REC

11.2.7 should be deleted, since the target never performs error recovery, and never sends an REC. (Question from Matt Wakeley)

This text is installed precisely because the target is intended to use REC to detect an unsuccessful transmission of FCP_CONF. Do we really want to make the change Matt proposes?

7.4 Change location of SRR

Annex B - This annex documents the “new” features of ABTS and REC. SRR needs to be added as well. (Question from Matt Wakeley)

Not accepted. See technical comment for 7.1

8.0 Other stuff to worry about

8.1 REC_TOV set/sense capability missing

At present, there is no mechanism to set or test the value of REC_TOV, but there is a mechanism to modify RR_TOV. Since there are a number of requirements that establish relationships among these values (for example, RR_TOV must be at least 3 times REC_TOV), we must either fix these values or provide a mechanism to detect and change them.

8.2 In-order delivery requirements

What fails if out-or-order fabrics are used? Can it be fixed? Is it a small enough set of the recovery mechanisms that we should continue to prohibit all out-of-order behavior, or do most of the mechanisms work correctly even with out-of-order delivery?

8.3 Is the discovery table any help?

What else should be included in the table? See FCP-2, rev 02, section 4.6)

8.4 Behavior of PRLI

There is an implicit assumption in the choice of bits in the PRLI request payload and in the PRLI accept payload that the PRLI request is always performed by an initiator. Since devices can label themselves as both and since there is no explicit rule that says the PRLI request is always done to a device that is only a target, I assume that the bits useful for initiators should be placed in both the PRLI request and the PRLI accept payload.

The following bits were copied over from table 9 to be placed in section 6.2.7, table 10.

Confirmed Completion Allowed

Data Overlay Allowed

I have not yet adjusted the text to clearly identify the bits as being sourced by initiators and not set by targets. The reason is that the PRLI image creation capabilities seem to be somewhat at odds with the informative and negotiative intent of the capabilities bits in FCP-2. This will be addressed as a separate issue in 8.5, which proposes that process associators be made obsolete in FCP-2.

8.5 Obsolete process associator

There is an informal proposal for making process associators obsolete, at least for FCP-2. I will be making that proposal formal for the next FC and FCP-2 meetings.

Note that the Process Associator definitions do not create a consistent architecture with SCSI and with PRLI. The problem is:

1) Process associators do not take part in the SCSI LUN or initiator addressing.

2) Process associators do not take part in separating CRN or exchange recovery.

3) The theoretical basis for process associators implies that independent processes are operating in the host. However, reservation protocols use as their primary parameter various initiator port identifiers, implying that the independent processes are not independent for at least that major part of the SCSI behavior.

4) PRLI has some problems separating initiator/target capabilities by process associator, since the process associator is not part of the addressing structure.

The best way to avoid having to figure out rational answers to all these questions (which will inevitably violate other standards) is to simply make them obsolete and not use them.

This is made more compelling by the fact that they are essentially unusable with their present definitions.

8.6 Name server extensions

Two recent documents have addressed FCP-2 specific capabilities in the Fibre Channel name server. One possible approach to these is to include them as a normative annex in FCP-2 as the first of several protocol specific extensions. A second possible approach is to include these as an FCP-2 device server definition within the body of FCP-2. I still personally prefer that the name server accept these in some format or other, possibly pre-documented by the appropriate protocol documents. No decision has been made on this subject yet.

Issues in this area include:

Possible solutions:

Name server extensions, as in 98-508v0.

Create an FC-4 specific object name space which is defined by the individual protocol document. (preferred at that particular meeting).

Creation of a new “device attribute server”, documented in a separate standard.

Do all this through non-standard mechanisms using CIM as the model of the attributes, possibly in a vendor specific manner.

Note that initiator attributes must also be defined in this way.

Charles Binford suggests that zoning access to this information from the viewpoint of the target may be critical. He points out that some vendors already provide separate images of Inquiry strings and logical unit lists on the basis of initiator. This can only be determined when asked for from the initiator that actually is interested.

8.7 REC extensions

Ed Gardner in e-mails of April 29 and May? suggested that REC could be improved by making it possible to obtain information across multiple exchanges. There has been no work and no proposal about this subject. I see no strong motivation to create this extension.

8.8 Device ID command

Bill Dallas in an e-mail on June 4 suggested that it would be interesting to create an INQUIRY-like command that would always be able to identify a device's unique id even while the device was executing self-test. There has been no work and no proposal about this subject. I believe it would not go in FCP-2 anyway, but probably belongs in SPC-2.