

March 31, 1999

T10/99-168 revision 1



To: T10 Membership
From: Bob Snively, Sun Microsystems
Subject: FCP-2 Changes from Revision 0 to Revision 1
This document reflects the following actions on FCP-2.

Discussion of FCP-2 in FC-TAPE and Fibre Channel meetings of December and February.

Resolution of FC-TAPE review comments for those sections that have been moved to FCP-2.

Approval of SAM-2, SPC-2, and other standards changes that change text in FCP-2, including:

NCITS T10/98-241r1 (Obsolete flag bit)

NCITS T10/99-158r1 (Long CDB definitions)

Identification of editorial inconsistencies resulting from the above modifications.

Resolution of technical problems identified outside the normal meetings. (These are clearly identified with an editor's note and square brackets to make location of these issues simple.)

The corrections and review actions are individually described below and the selected resolutions clarified where necessary.

1.0 Technical Changes

1.1 Capture FC-TAPE comment resolutions in clause 11

FC-TAPE comments addressed a significant number of irregularities in the chapter about link error recovery, clause 11. These comments are reflected in those passages marked with change bars. The changes have not been exhaustively verified.

1.2 Annex C completed

Annex C did not previously have any contents due to program incompatibilities with the files for FC-TAPE. The annex is now completed. As the result of FC-TAPE comment resolutions, figures C.7, C.9, C.10, and C.12 have been modified. The FCP_TOV timer has been replaced with the REC_TOV timer.

1.3 Make flag bit obsolete

NCITS T10 approved in its March meeting proposal 98-241r1 to make the flag bit obsolete in the CDB control byte. This proposal had already been included in FCP-2 revision 0 on the assumption that it would pass.

1.4 RR_TOV timer

Jim Coomes references a PLDA annex (Annex D) that contains a parameter for the Fibre Channel control mode page (page code 19h) allowing the definition of a resource recovery timer in a mail of Feb 25, 1997. This has been approved by the working group. It 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 have been selected by phone discussions with key contributors. 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	1	0	0.001 seconds
0	1	1	0.1 seconds
1	0	0	10 seconds

Table x: Values for RR_TOV units

Byte 6			Units of measure for RR_TOV
bit 2	bit 1	bit 0	
Other			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. Those functions are specified by FC-PLDA and by section TBD of this standard.

The supplementary text for this timer has not yet been created in FCP-2. See 3.1 of this letter.

1.5 Long CDB

Gerry Houlder has proposed an extended CDB capability. This capability was defined for SPI-3 and for FCP-2 by the approved document T10/99-158r1. The table from the referenced document replaces the table defining FCP_CMND in section 8.1 of FCP-2.

There is an editorial mistake in NCITS 99-158 revision 2. FCP-2 requires that all task management functions, including CLEAR ACA, are individually carried in a separate FCP_CMND payload. Document 158 incorrectly specifies that CLEAR ACA may be carried in the task management function with a valid CDB. The text introduced into FCP-2 was modified to correct this. Revision 3 of 99-158 corrects this error and applies the same change to SPI-3.

1.6 Resolution of the “Tape problem”

FC-TAPE presently has a number of specialized error recovery ELS’s proposed for implementation of the SCSI management in both class-2 and class-3. The standardized use of this is defined in FCP-2 and referenced in FC-TAPE.

The ELS’s defined for tape recovery will be placed in an FCP-2 annex until they can be included in FC-FS. The recovery procedures will be specified in a new section of the FCP document body. (Done.)

The method for using these is presently specified in section 11. This section needs a lot of work still, since it is stolen intact from FC-TAPE and must be modified to have the form of a standard instead of the form of a profile. (Done, but lots of editorial work is still required.)

Examples of typical error recovery events have been placed in Annex C. The information is based on the old FC-TAPE Annex B.

1.7 Clarification of ABORT TASK function

At present, all task management functions except ABORT TASK are marked as complete by an FCP_RSP IU. ABORT TASK in SAM-2 has the interesting property of being acknowledged by the device server, but removing the acknowledgment that would normally have been presented by the aborted task. FCP has always been a bit unclear about distinguishing between the ABORT TASK function and the Recovery Abort function. This needs to be clarified and perhaps corrected. Gen-Hwa Chiang's mail of 5 Nov 1996 pointed out this lack of clarity. Dave Peterson has provided additional proposals about this issue.

The first part of this clarification requires section 8.1.2.1 to indicate that ABORT TASK is performed by executing a recovery abort function.

The second part of this requires me to analyze the error management functions newly installed and compare them with the recovery abort capabilities and other text. This may require additional changes in both sets of text. This is not yet done.

1.8 Data overlay description modifications

In section 8.3, the data overlay text suggested by FC-TAPE comment 30 from Sun Microsystems is installed as follows:

If the amount of data returned does not match FCP_DL, the error detection and recovery procedure described in clause 11 may be invoked or the FCP I/O Operation may be terminated with a recovery abort or other failure indication. The manner in which a SCSI Initiator determines that the correct amount of data is returned is outside the scope of this technical report. Data that has been retransmitted and overlaid shall be counted only once.

1.9 Data overlay usage restrictions

FC-TAPE has identified a number of cases that may require restrictions on the data overlay capability during retry. Comments about these cases were addressed in the FC-TAPE meetings and the appropriate solutions placed in the chapter which has become clause 11 of the FCP-2 document. These changes include:

11.2.5: FCP_XFER_RDY for write retries shall not specify an RO smaller than that specified in the SRR.

11.2.6: If an initiator does not have enough information to create a correct RO in the SRR ELS, it will set the RO to zero and retry the whole data transfer. This may be the case especially when data had been transmitted out of order.

Note that some of the data overlay usage restrictions are not FCP restrictions, but rather are additional restrictions imposed by FC-TAPE that are suitable for sequential devices, but not block devices. See particularly sections 5.2 and 5.3 of this document.

1.10 FCP_RSP field length

There is incomplete documentation of the required minimum content of the FCP_RSP field. While the FCP_SNS_LEN and FCP_RSP_LEN have notations indicating that their presence is mandatory even if the field is unused, the FCP_RESID has no such notation. Section 8.4.7 is updated to reflect this.

1.11 Expected corrections to clearing table

Dave Peterson prepared updated revisions of the clearing table (see section 4.6) which include corrections resulting from FC_TAPE document review. Those have been installed.

1.12 SRR definition

At present, SRR is treated as an Extended LINK_Data frame, to be moved to FC-FS at some time in the future. FC-TAPE comment 56 by Robert Kember indicates that, since there are SCSI specific parameters in it, it might be better to document the SRR as an FC-4 link service in the FCP-2.

The R_CTL field of SRR will be changed to 32/33H and documented in clause 7 of FCP-2 instead of in Annex B.

Note that many FC-TAPE corrections have been installed in the ELS definitions.

1.13 FC-TAPE corrections to annex B

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.

2.0 Review Actions requested by committee

2.1 Review section 9.1.3 with respect to T10/98-106r1

The text in FCP-2 has been modified to fit the requirements for a standard, including use of the active voice where possible and use of the key words "shall", and "may". References to the proper standards are provided in FCP-2 where such allusions were made in the proposal document. Some tutorial text outside the scope of the standard is deleted, especially when describing the DLM bit. In addition, those bits that act differently in the presence of loops and point-to-point links are identified and the differences clarified.



An error was inadvertently introduced into FCP-2 in the definition of the DTOLI bit. FC-TAPE and T10/98-106r1 allow the target to generate an error LIP if the link fails, regardless of the state of the DTOLI bit. FCP-2 only allowed the error LIP to be generated if the DTOLI bit was zero. FCP-2 is corrected. Note that this correction gives the target the curious property of introducing a LIP into the link if the link provides the appearance of a failing link during the very initialization when it was not supposed to provide a LIP.

An error was inadvertently introduced into FCP-2 in the definition of the DLM bit. If the DLM bit is zero, FCP-2 indicates that the target shall participate in loop master arbitration and may become master. Both FC-TAPE and T10-98-106r1 indicate that the target may become master, but do not require it to participate in loop master arbitration to become master. FCP-2 will be corrected.

2.2 Review 9.1.2, 9.1.3 with respect to FC-TAPE Annex E

FC-TAPE Annex E was a temporary repository for the information contained in section 9 of FCP-2. Significant updates have occurred since Annex E was written. The Fibre Channel Control pages have been divided to separate those associated with LUNs and those associated with targets. In addition, enabling bits for new FCP-2 functions have been added. The FCP-2 text significantly clarifies the text of Annex E with respect to the FCP-2 functions, but makes no technical change. Changes similar to 2.2 above in the FCP-2 text were required to make the language of the document conform with the T10 practices. Both inadvertent errors found in 2.2 above were also present in this comparison, and FCP-2 was changed to correct them.

There do not appear to be any other technical differences between the bit definitions of FC-TAPE Annex E and FCP-2.

3.0 Technical changes to be approved

3.1 Timer definitions

Previously, the definitions of the timers used by clause 11 in the error recovery process were all outside this document. There are tables in section 8 of FC-TAPE, corresponding tables in FC-PLDA, and a number of definitions scattered among the FC-PH and FC-AL documents. I have moved the timer information into a new section of FCP-2, together with references to the defining standard (not profile) if applicable. For those without a defining standard, FCP-2 should become the defining standard.

A first draft attempt has been made to change this section from the format of a profile to the the format of a standard. Most timers are defined by reference

to the relevant document. RR_TOV, REC_TOV, and ULP_TOV are defined as FCP-2 timers.

YES NO ABSTAIN

3.2 REC_TOV timer implementation

The present FC-TAPE and FCP-2 documents indicate that there shall be one REC_TOV timer for each exchange. This must include all exchanges, not just the ones that are active at any given instant, and the number may be very large. I would propose that we modify this requirement such that a single timer can be used and referenced for all exchanges, making the first REC access executed for a particular exchange be between 1 and 2 REC_TOV after the command is sent. That has the possibility of simplifying the implementations significantly, keeping only a single bit per exchange. No change is made yet.

YES NO ABSTAIN

3.3 Negotiation of FCP Recovery capabilities

There have been several discussions about how the more sophisticated functions of FCP, including execution of REC, execution of SRR, support of FCP_CONF, and support of CRN are discovered and enabled.

The committee has elected to put FCP_CONF in the PRLI. That is presently installed.

The committee has elected to put CRN in the MODE SELECT page for Fibre Channel LU control. That is presently installed.

There has been no solution presented for SRR and REC. At present, the way you find out is to try an REC or an SRR and quit if you get LS_RJTs. I believe this is an adequate mechanism and I have not included or invented any supplementary discovery or negotiation mechanism.

YES NO ABSTAIN

3.4 Obsolete mixed command/data and data/response

At present, there are no implementations that allow combining the command and data in a single IU nor any implementations that allow combining the data and response in a single IU. All implementations use separate IUs (and therefore separate sequences) for all individual phases. Of course, the phases may be transferred during the same loop tenancy for improved FC-AL efficiency. I propose that the FCP-2 be simplified by removing and making obsolete the use of mixed command/data and mixed data/response. That would eliminate some text in section 8.1, remove the necessity to clarify sections 8.3 and 8.4, eliminate the PRLI parameters of command/data mixed allowed and data/response mixed allowed, and remove IUs T8, T9, T10, T11, I6, and I7. It would also save the effort of resolving and documenting the error recovery procedures for these cases.

There has been considerable discussion of this subject on the reflect, and the consensus has clearly been in favor of removing the function, but some concerns remain that must be addressed during the study of this proposal. No change has been made and further study is required, but I believe this is a sound and useful proposal.

YES NO ABSTAIN

3.5 Obsolete First Burst Length parameter definition

In section 9.1.1, the parameter First Burst Length in the Disconnect/Reconnect mode page is defined in the following manner.

The first burst size field indicates the maximum amount of data that a target may transfer for a command during the same interconnect tenancy in which it receives the command. This value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates that there is no first burst size limit. This value shall be implemented by all FCP-2 devices. The application client and device server may use the value of this parameter to adjust internal maximum buffering requirements.

This definition was correctly criticized by George Penokie who preferred that the definition be: "The first burst size field indicates the maximum amount of data that an initiator may send as a first burst data transfer."



While George's criticism is valid, there are additional errors in the definition. Since these errors have never been found, it is my first inclination to make this parameter obsolete. If the committee feels that this is not an appropriate solution, then the following text should resolve or identify most of the problems associated with this definition:

The First Burst Size field indicates the maximum amount of data that can be transferred in the first FCP_DL IU for an FCP I/O operation. This value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates that there is no first burst size limit. This value shall be implemented by all FCP-2 devices. The application client and device server may use the value of this parameter to adjust internal maximum buffering requirements in the following manner.

The value of the First Burst Size field may be assumed by initiator and target. If the First Burst Size value is set by the initiator using a MODE SELECT command, the actual value accepted by the device should be verified using the MODE SENSE command.

If Write XFER_RDY Disabled is set to 1, the initiator shall manage write data transfers such that the first FCP_DATA IU does not exceed the size specified by the First Burst Size field. Initiators using FC-AL or FC-AL-2 links shall not transmit more than the first FCP_DATA IU in the same tenancy with the FCP_CMND IU.

If Write XFER_RDY Disabled is set to 0, the target shall manage both read and write data transfers such that the first FCP_DATA IU transmitted by the target or requested by the target does not exceed the size specified by the First Burst Size field.

A similar concern from Ralph Weber, which tied First Burst Size to transfers within a tenancy, would also be corrected by these changes.

No change has been made in FCP-2 to address this question yet.

YES NO ABSTAIN

3.6 Interpretation of FCP_RSP

As a result of review comments received by Dave Peterson on the FC-TAPE document, he has requested that the following test be placed at the proper location in FCP-2.

The RSP_CODE is independent of the SCSI Status and should be examined prior to interpretation of the SCSI Status.

After reviewing the section, I believe it is necessary to clarify a number of the relationships among FCP_RSP fields. The following changes were made:

Section 8.4: The following text was installed to indicate that non-zero values in bytes 10 and/or 11 require that the initiator examine additional information to determine whether the FCP I/O Operation completed correctly.

A non-zero value in either byte 10 or byte 11 should be examined by the application client to determine whether a failure, a retrievable temporary condition, or an expected response occurred.

Sections 8.4.2 through 8.4.5: Text was installed to indicate that a “1”b in any of these flags required the application client to examine the parameters pointed to by the flag to determine if meaningful information was available in the referenced fields.

Section 8.4.9 and 8.4.10: At present, an FCP_RSP_LEN of zero with the FCP_RSP_LEN_VALID bit equal to one is a valid configuration. I believe that this was not intended and have changed the document to require no FCP_RSP_INFO field to have an FCP_RSP_LEN_VALID of 0 and to prohibit a zero length if the FCP_RSP_LEN_VALID is 1.

Section 8.4.11: The FCP_SNS_INFO shall be zero length and the FCP_SNS_LEN_VALID bit shall be 0 if there is no sense information to be presented. FCP devices shall use autosense.

The above changes have been included in revision 1.0 of FCP-2.

YES NO ABSTAIN

3.7 SRR Reason Code incorrect

Dave Peterson has reviewed the SRR LS_RJT reason code, 00052A00h and found that it is mislabeled. The code should be labeled:

Logical Busy/Unable to supply requested data.

The text of section 7.1 is updated accordingly. **This may require further verification.**

YES NO ABSTAIN

3.8 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. I believe that the first paragraph should be rewritten to reflect the SCSI-specific use of SRR as an initiator-only function.

YES NO ABSTAIN

3.9 Make class independent recovery optional for FCP-2

Section 11 skates around this issue by suggesting that FCP-2 link recovery is tailored for certain types of devices, specifically tapes. The following sentence replaces that and makes it clear that class independent recovery as defined by section 11 is an FCP-2 option.

The class independent FCP-2 link error recovery procedures documented by this clause are optional and need not be implemented by FCP-2 compliant devices. Exchange level recovery is acceptable for many FCP-2 applications. The link error recovery procedures described here are especially useful for those devices which may maintain order-dependent implicit state information and have a long or complex exchange level recovery mechanism.

YES NO ABSTAIN

4.0 Editorial Changes

4.1 Numerous small corrections

Cross references, spelling corrections, grammar corrections, and editorial style corrections were made in many areas. None have technical meaning.



4.2 Parameters in small caps

All sections:

Parameters should be identified by small caps instead of full-size caps. This correction will be installed in the next revision.

4.3 Editorial improvement of FCP_CMND payload table

The change in the FCP_CMND payload included in 158r2 proposed a modified table format, more in keeping with SPI-3 and other related standards. This format was adopted for the information units in chapter 8 and makes clearer the relationship among bits and bytes in those information units. The FCP_CMND was reorganized, but no text was changed except to remove the F_CTL definition, which was a useless construct.

The same improved table structure is applied to FCP_XFER_RDY, with no additional text changes required.

The same improved table structure is applied to FCP_RSP, with some additional text modifications to remove the sub-grouping of fields called FCP_STATUS. This also required minor editorial changes to section 8.3.

4.4 REC_TOV is the new name for FCP_TOV

This change was made prior to moving the corrected versions of clause 11 from FC-TAPE to FCP-2.

4.5 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.

4.6 Correct description of FCP_CMND/FCP_DATA

Giles Frazier, in reviewing section 8 of FCP-2 revision 0 found the following error. Since the last sentence of section 8.1 starts with the phrase "in all other cases," it implies that unless BOTH XFER_RDY IUs are disabled AND Command/Data Mixed is allowed, then the FCP_CMND IU must transfer sequence initiative.

Actually, if XFER_RDY IUs are disabled but Command/Data Mixed is NOT allowed, then the FCP_CMND IU for a write should NOT transfer sequence



initiative. Instead, the FCP_CMND IU should be immediately followed by the first or only FCP_Data burst. (This is required by sections 6.2.6.13 and 8.2.)

After careful review, this paragraph appears to be redundant with section 6.2.6.11. The text is now deleted from section 8.1.

4.7 Include FC-AL and FC-AL-2 references

FC-AL and FC-AL-2 were added to the references in sections 2.1 and 2.2 respectively.

4.8 Correct spelling of PLPB

In section 9.1.3, the mnemonic PLPB was incorrectly spelled as PLB. The document was corrected.

4.9 Correct reference to SAM-2

The reference to SPC-2 in section 8.1.1.4, LOGICAL UNIT RESET paragraph, is corrected to be a reference to SAM-2.

4.10 Editorial corrections on Disconnect-Reconnect page

The word “interconnect latency” is corrected to be “interconnect tenancy”.

4.11 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.

This text is not yet installed in the document.

4.12 Section 6.3.2.1, reference needs correction.

The references to annex A in section 6. were corrected.

4.13 Annex B modified to be a standard, not a profile

A number of editorial modifications were made in Annex B to more closely reflect the style of a standard. Sections B.1 and B.1.1 were modified. The LS_RJT code name was corrected to “logical busy”.

5.0 Items not planned for inclusion

5.1 Definition of mode page parameters for FCP, DTDC field

PLDA specifies in table 18 a set of required and prohibited mode page fields for directly attached devices. With one exception, these are consistent with FCP-2. Approved document 95-348r1 specifies that the DTDC and Dimm fields are not appropriate for FCP-2 devices and shall be zero. PLDA indicates that a DTDC value of 011 is allowed. This value rather meaninglessly requires that a target shall transfer all data for a command and complete the command within a single interconnect tenancy. I believe that 95-348r1 is correct and PLDA is incorrect and have taken this approach in FCP-2. This will also be reviewed by members of the committee. At least one implementation may have used these bits for forcing FCP_RSP and FCP_DATA into separate tenancies. The implementor is reviewing this work to determine if and how it should be included.

No proposal has been presented and no change will be made.

5.2 Restrict data overlay to error recovery

As a result of FC-TAPE comment 13 from Charles Binford, text was placed in the FC-TAPE document that constrained the usage of data overlay to error recovery cases only. As a result, the following text has been placed in section 9.2.2 of FC-TAPE.

SCSI Targets conforming to this profile shall be capable of performing link level error recovery using data overlay if allowed by the SCSI Initiator. The use of data overlay shall only be allowed in response to an SRR (i.e. error recovery).

This restriction is not required by FCP and is not required for proper execution of error recovery if both initiator and target have the necessary capabilities. No change will be made.



5.3 Continuously increasing sequence count

The solution to FC-TAPE comment 13 from Charles Binford incidentally required continuously increasing sequence count for class 3 streamed reads, including the FCP_RSP sequence. I believe this is covered adequately by the appropriate FC-PH documents and need not be repeated in FCP-2, although of course it is an appropriate item for FC-TAPE to profile.

No change will be made.

5.4 Data checking when doing overlay or out of order xfer

Charles Binford's FC-TAPE comment 37 addresses the verification of complete data transfers when doing data overlay. At present, FCP-2 indicates that this is an implementation detail outside the scope of the standard in section 8.3 in the following manner:

If the amount of data returned does not match FCP_DL the error detection and recovery procedur described in clause 11. may be invoked or the FCP I/O Operation may be terminated with a recovery abort or other failure indication. The manner in which a SCSI Initiator determines that the correct amount of data is returned is outside the scope of this standard. Data that has been retransmitted and overlaid shall be counted only once.

No additional change will be made.

5.5 FCP_RSP IU length

Some older devices have elected to provide only 12 bytes of the FCP_RSP IU if there are no valid contents in the FCP_RESID, FCP_SNS_LEN, and FCP_RSP_LEN. This is contrary to the present document, which requires that those fields always be included, implying a standard length of 24 bytes for FCP_RSP unless FCP_RSP_INFO or FCP_SNS_INFO are present. A control bit in the Fibre Channel LUN control page has been informally discussed to allow 12 byte FCP_RSP fields to be presented. No formal proposal has been presented. No change is being made.



5.6 Automatic persistent reservation registration update

The FC-TAPE comment 56 by Robert Snively included a number of issues. One requested that the persistent reservation registration information and reservation information be adjusted during port discovery so that reservations would remain fixed to a particular host world-wide name even if D_ID/S_ID values of the initiator were changed. That had already been installed in section 5.3 of FCP-2. No further change is required.

5.7 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. **There is still an editor's note in this paragraph.**

5.8 FCP_CONF and linked commands

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

5.9 Obsolete Third Party Process Logout

At present, this has been discussed, but no firm conclusion has been reached about whether it is appropriate to remove TPRLO. There has been some use of it for breaking the authority of members of a cluster to access certain devices. This can be done more effectively using Persistent Reservations.



Further study is required.

5.10 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.

5.11 SRR Transfer of initiative

The FC-TAPE comment 55 by Bob Kembel suggests that the transfer of initiative forced by the SRR command is outside the normal transfer of initiative rules. No change in the relevant text is proposed.

5.12 Retry offset

Rob Basham of IBM has presented discussions in the committee meetings that indicate that SRR retries of fixed block reads could create problems on some tape device implementations. The FCP-2 document allows retries to begin at the beginning of the SCSI data transfer to accommodate such implementations. No change in the text is proposed.

5.13 Review Enable Automatic FCP_RSP generation

Automatic response generation allows the generation of the FCP_RSP in the same tenancy as the last inbound or outbound data transfer. A control bit has been informally suggested in the Fibre Channel Port Control page, byte 4,



bit 0. No proposal has been presented about this subject and no changes are made in FCP-2 revision 1.

5.14 Correction of DDIS bit

The disable discovery function is presently described as allowing a device to operate normally even though no discovery ELSs (ADISC or PDISC) have been exchanged after initialization. This allows any tasks being executed to continue as soon as the initialization process is completed.

The question was raised about the behavior of login in the presence of the DDIS bit. I believe that DDIS has no effect on the login process and requirements, since that functionality is managed by the ALWLI (Allow login without initialization) bit. No formal proposal has been made and no changes are made in FCP-2 revision 1.

5.15 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.

5.16 CRN recovery

At the committee meetings, there was some informal discussion of what happens when a CRN violation occurs because a command is unsuccessfully recovered. If first level recovery fails and second level recovery is also unsuccessful, the second level recovery terminates with an explicit or implicit logout. While this needs to be properly verified, I believe this solution is complete.

