

September 30, 1999

T10/99-247 revision 1



To: T10 Membership
From: Bob Snively, Sun Microsystems
Subject: FCP-2 changes from revision 02 to revision 03.

At the FC-TAPE/FCP-2 meetings of July 13, 1999, August 4, 1999, and September 14, 1999, document T10/99-211 revision 0 and FCP-2 revision 02 were extensively considered. The first part of this document concentrates on those items from document 99-211 that were the subject of controversy or discussion in either of those meetings or by E-mail. In some cases, these items had not yet been included in FCP-2 because resolution was uncertain. Where applicable, final conclusions are identified.

Additional discussion items were brought forward in those meetings. Those items were discussed and, in most cases, resolved. The second part of this document identifies the issues and the resolutions for each of those comments. The items include:

Charles Binford, "Comments on FCP-2 v2" dated 8/4/99.

Stewart Wyatt, e-mail "FCP Revision 2 Comments" dated Sept 9, 1999

1.0 Approval of changes, FCP-2 rev 01 to 02

The following changes that were included in FCP-2 revision 2 were the subject of discussion at the July and August meetings. Where there was no controversy, the subjects are deleted from this document. Those that were discussed and approved without modification are listed below..

Technical changes discussed and approved in FCP-2 revision 2

Section	Ref	E/T	Summary	Approved
7.1	1.1.1	T	SRR should be moved	7/13/99
7.1	1.1.2	T	Correct reason code	7/13/99
7.1	1.1.3	T	Correct reason code	7/13/99
7.1	1.1.4	T	Verify reason codes	7/13/99
9.1.1	1.1.5	T	Clarification of Buffer Empty Ratio	7/13/99

Discussion items with approval not to install

Section	Ref	E/T	Summary	Approve
10	2.1.3	T	Default E_D_TOV	7/13/99
11.2.7	2.4.2	E	Target never sends REC	7/13/99

1.1 Technical corrections discussed and approved unchanged

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. They were then approved at the July meeting.

1.1.1 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. This was approved in the July meeting.

[\[done in rev 02\]](#)

1.1.2 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. This was approved in the July meeting.

The actual value of the “unable to perform the command request” reject code is 09. That value is used in revision 03 of FCP-2.

[\[done in rev 02\]](#)

1.1.3 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. This was approved in the July meeting.

[\[done in rev 02\]](#)

1.1.4 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. This was approved in the July meeting. The reference was changed to FC-FS.

[\[done in rev 02\]](#)

1.1.5 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. This was approved in the July meeting.

[done in rev 02]

2.0 Requested changes installed, FCP-2 rev 02 to rev 03

The following changes have been placed in FCP-2 between revision 02 and revision 03. These changes have been approved in those technical meetings since the publication of revision 02. 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. Typographical and spelling errors are corrected without comment. All changes in FCP-2 except typographical and spelling errors are marked with change bars.

Technical changes installed in FCP-2 revision 3

Section	Ref	E/T	Summary	Approved
8.1.1.4	2.1.1	T	Mandatory and optional task management	7/13/99
TBD	2.1.2	T	Rules for ELS generation before Login	
10.0	2.1.3	T	Default E_D_TOV	9/14/99

Editorial changes installed in FCP-2 revision 3

Section	Ref	E/T	Summary	Approve
All	2.2.1	E	Parameters in small caps	NA
11	2.2.2	E	References for error examples	NA

Editorial and Technical changes from LSI comments

Section	Ref	E/T	Summary	Approve
4.2	3.1	E	Editorial	
4.2	3.2	E	IU termination	

Editorial and Technical changes from LSI comments

Section	Ref	E/T	Summary	Approve
4.4	3.3	E	Clarification of retransmission	
4.4	3.4	T	Restriction of FCP_CONF usage	
table 2	3.5	E	Clarify generation of REC	
3	3.6	T	Include LUN Reset in table	
4.8	3.7	T	Recovery only for FCP sequences	
6.1	3.8	E	Editorial	
6.2.6.	3.9	T	Confirmed Completion Allowed bit	
6.2.6.9	3.10	T	Data overlay required for retry	
7.1	3.11	T	Reject of retransmission requests	
7.1	3.12	T	XFER_RDY retry	
9.1.1.3	3.13	T	Bus Inactivity Limit clarification	
9.1.1.4	3.14	T	Disconnect Time Limit clarification	
9.1.1.6.	3.15	E	Burst size during write clarification	
9.1.1.10	3.16	T	First Burst Size clarification	
Table 28	3.17	T	Choice of RR_TOV default	
Table 28	3.18	T	Clarify FLOGI	
11.1.1, 11.1.2	3.19	T	Class 1 and 4 treatment	
11.1.2	3.20	T	Missing ACK after FCP_CONF	
11.1.2	3.21	T	ACK(abort) to target	
11.2.1	3.22	E	Editorial	
11.2.1	3.23	E	Editorial	
11.2.3	3.24	E	Read XFER_RDY obsolete	
11.2.3	3.25	T	Permission to do recovery	
11.2.4	3.26	T	REC data count	
11.2.4	3.27	T	Recovery of 0 Bytes Xferred case	
11.2.4	3.28	T	Hold exchange information	
11.2.5	3.29	T	RO during recovery	

Editorial and Technical changes from LSI comments

Section	Ref	E/T	Summary	Approve
11.2.6	3.30	T	FCP_RSP retransmission	
11.2.6	3.31	T	Recovery RO	
11.2.7	3.32	T	OX_ID reuse and RX_ID	
B.3.1	3.33	T	OX_ID and RX_ID in REC	

I

Editorial and Technical changes from HP comments

Section	Ref	E/T	Summary	Approve
3.1	4.1	T	New glossary terms	
3.1.16	4.2	E	Reference for FQXID	
4.1	4.3	E	References updated for documents	
4.6	4.4	E	editorial	
4.9	4.5	T	Clarify implicit PRLI/PRLO beyond scope	
6.2.6.5	4.6	E	editorial	
6.2.6.9	4.7	E	editorial	
6.2.6.10	4.8	E	editorial	
11.1.2	4.9	E	Correct directional reference	
11.2.7	4.10	T	Target retry of FCP_CONF	
A.4	4.11	E	editorial	
C	4.12	E	Correct references and table C.1	
D	4.13	E	Correct references	
D	4.14	E	Correct figure D.9	

Discussion items not yet included in FCP-2 revision 3

Section	Ref	E/T	Summary	Approve
TBD	2.3.1	T	Consideration of additional ELSs	
TBD	2.3.2	T	Name server extension	
6.2.7	6.1	T	Behavior of PRLI	
many	6.2	T	Obsolete Process Associator	

Discussion items rejected for FCP-2 revision 3

Section	Ref	E/T	Summary	Rejected
5.2, table 6	2.4.1	T	Treatment of redundant name spaces	9/14/99
11.2.7	2.4.2		Target never sends REC	

2.1 Technical changes for FCP-2 revision 3

2.1.1 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	

These choices were approved in the July meeting.

[\[done in rev 03\]](#)

2.1.2 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 is partially included in section 11.10, but more work remains.

This work item is not yet complete and should be included in FCP-2, revision 04.

2.1.3 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.

This value is defined as the default in FC-PH. See section 23.1.1.2 of document 99-145v0, the integrated FC-FS document.

After further discussion, Dave Baldwin will carry a formal proposal to the FC-FS working group to make the default value 2 seconds.

In FCP-2, the FC-FS reference for the default value will be provided and the present default specified. In addition, the methods for changing the value from the default value will be clarified.

[\[done in rev 03\]](#)

2.2 Editorial corrections approved at previous meetings

2.2.1 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.

[done](#)

2.2.2 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 04.**

2.2.3 MCM additions

Additional changes have been proposed for MCM management in the Fibre Channel Control Page in FCP-2 in document 99-206r2. This document was marginally approved within the SCSI plenary, so its technical content is included in FCP-2 revision 3. Note that there are a number of holes and incompletenesses in this work at present, including a clear requirement for the necessity of such a function.

[done](#)

2.2.4 DSA bit clarifications

A proposal has been put in place changing the name and clarifying the description of the DSA bit. The proposal, 99-226r2 from Jim Coomes, was approved in the last working group meeting. The changes are now installed in FCP-2 revision 3.

[done](#)

2.3 Discussion items not yet installed in FCP-2, revision 03

2.3.1 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.

The new text prepared by Dave Peterson is now included in annex E of the document.

In addition, this change has revealed that the conventions for numbers are not defined in section 3. The appropriate text modeled by SPC-2 is used.

[\[done in rev 03\]](#)

2.3.2 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.

After considerable discussion, the following approach has been selected.

Create an FC-4 specific object name space which is defined by the individual protocol document.

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. The meeting group determined that this was outside the scope of FCP-2.

A proposal document is required for this function.

2.4 Technical proposals rejected from FCP-2 revision 03.

2.4.1 LSI 008 * Treatment of redundant name spaces (Technical)**

page 14, section 5.2, table 6 The FCP_Port Identifier is insufficient for typical FCP installations that have redundant loops or redundant fabrics (and redundant FCP_Port Identifier address spaces). Wouldn't it be wonderful if this field was the WW PortName instead? If we remove Process Associators, then we could use all 8 bytes.

This violates the SAM architecture. This comment was rejected.

no change required

2.4.2 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 SRR or FCP_CONF. Do we really want to make the change Matt proposes?

As of the July meeting, the behavior as documented is accepted. More discussion may occur.

no change required

2.4.3 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.

Discussion items in the July meeting included:

- 1) The first REC should be able to come at any time without a minimum wait. This is apparently already allowed by the text.

- 2) REC_TOV has defined relationships with E_D_TOV, as does RR_TOV. As a result, it is probably acceptable to allow the limits on REC_TOV to be derived from the specified values of E_D_TOV and RR_TOV.

no change required

2.4.4 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.

This change was discussed in the July meeting. My notes indicate that the OX_ID and RX_ID must be examined together to identify an exchange. A scan should be done to find those cases where one alone was incorrectly referenced. The minutes indicate that this text was to be changed back to give no guidance, but my notes indicate that this suggestion applied only to the next problem listed in 99-211, not this problem.

This has been further modified by the realization that FCP_CONF will not work properly unless RX_IDs are used (see 3.32). As a result, the requirement is changed to require the use of RX_ID for class-2 targets.

done

2.4.5 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.

This change was further discussed in the July meeting. While the change is basically correct, it does not allow the special case of an ABTS sent to abort an exchange for which the RX_ID has not yet been received. The discussion ignored the subsequent paragraphs which indicated that RX_ID was a wild card if it had a value of "FFFF"X. Including this information, I believe the original change included in revision 02 is correct and have made no change to revision 03.

no change required

3.0 Consideration of LSI comments for inclusion in rev 3

Issues labeled as technical were the subject of discussion at the August meeting. Note that other comments may have technical content, but the content was considered to be merely a clarification or an obvious correction.

3.1 LSI 001 Editorial

page 7, section 4.2, 2nd paragraph Appears to be an extra 's' or missing characters in the phrase "including some command control s, addressing information,"

The typo "control s" is corrected to "controls".

done

3.2 LSI 002 IU termination

page 7, section 4.2, 4th paragraph The command IU terminates only if FCP_CONF was not requested.

The words:

The command status IU terminates the command.

are replaced with:

The command status IU indicates completion of the SCSI command. If no command linking, error recovery, or confirmed completion is requested, the command status IU also ends the exchange.

done

3.3 LSI 003 Clarification of retransmission

page 9, section 4.4, 2 nd paragraph In the last sentence, we need to clarify what information the target may be requested to retransmit - i.e. the FCP_RSP data.

The word "information" is changed to "FCP_RSP".

done

3.4 LSI 004 * Restriction of FCP_CONF usage (Technical)**

page 9, section 4.4, 4th paragraph Last sentence restricts FCP_CONF from being used for non-queued, non-check FCP_RSPs. Why? This sounds like a profile type restriction, not a standard.

This was discussed in the August meeting. Dal Allan provided a model for use in defining and refining the usage of FCP_CONF. The model had to be extended to allow the described behavior.

done

3.5 LSI 005 Clarify generation of REC

page 10, table 2 The / notation in the Initiator performs REC is confusing notation. Use a note or some other documentation method to convey information.

The text was changed to read: None required, Process Login allowed.

done

3.6 LSI 006 ** Include LUN Reset in table (Technical)

page 11, table 3 Should LUN Reset be in this table?

LUN Reset is required and was installed.

done

3.7 LSI 007 ** Recovery only for FCP sequences (Technical)

page 11, section 4.8 Intro to table need to add words indicating references to sequences and exchanges are FCP sequences and FCP exchanges ELS s and other PH exchanges follow PH rules.

The appropriate text will be added.

done

3.8 LSI 009 Editorial

page 20, section 6.1 I no longer consider PRLI to be a new extended link service. I suggest dropping the word new.

The change was made.

done

3.9 LSI 010 ** Confirmed Completion Allowed bit (Technical)

page 24, section 6.2.6.8 The title of the bit Confirmed Completion Allowed implies this is an initiator capability only (and I agree). However, the wording of the text implies the bit also indicates whether or not the target will ever ask for it (information I don't think is required). Specifically I'm referring to the phrases indicating that it supports . . . for both its target function and. . . and If the responder does not support. . . . It seems to me that wording closer to the Data Overlay Allowed text would be more appropriate (or rename the bit).

After the requested review, the editor agrees with Charles Binford. The text is modified to be similar in content to the text of the Data Overlay Allowed bit.

done

3.10 LSI 011 * Data overlay required for retry (Technical)

page 24, 6.2.6.9 Data Overlay Allowed I d suggest that this bit be required (shall be 1) if the new Retry bit is set.

The suggested change is accepted.

done

3.11 LSI 012 *** Reject of retransmission requests (Technical)

page 30, 7.1, 1 st paragraph at top of page Why is a target prohibited from rejecting retransmission requests for XFER_RDY or RCP_RSP. This seems an unreasonable requirement. It is not obvious to me that tape drive (for example) can in all cases successfully recover an interrupted write command. Further, FCP_RSP may not be available if a device is a bridge and serving both disk and tape. The device may support SRR, yet the LUN with the error may be a disk.

If this comment is rejected, then the flavor of the XFER_RDY (i.e. Read or Write XFER_RDY) needs to be clarified (specify Write XFER_RDY).

The general principle that should be followed is that retransmission requests should always be accepted unless there is an error or other condition that prevents their execution. A device should not commit itself to retry, then reject all possible attempts to perform a retry. In view of this, the suggested modification is probably desirable.

done

Additional issues concerning clarification of the valid responses to SRR were raised in the discussion. These are clarified in the appropriate sections.

1) The target may choose to end a command with FCP_RSP rather than retry.

included in 11.2.6, rev 02

2) SRR ACC timing with respect to the response is not defined. This should be the same value specified for normal ELS responses and need not be stated. **Not included.**

3) FCP_RSP may not be retryable on some LUNs of a device.

done

3.12 LSI 013 * XFER_RDY retry (Technical)

page 29, 7.1 The SRR description does not comment on whether or not a device supporting Read XFER_RDY should retransmit the XFER_RDY when is has been requested to retransmit Read data. This scenario needs to be clarified.

Read XFER_RDY is made obsolete. See 3.24. In any case, this information is described in chapter 11.

no change for this comment

3.13 LSI 014 ** Bus Inactivity Limit clarification (Technical)

page 48, 9.1.1.3 Bus inactivity limit The definition seems inconsistent. - 0xffff indicates that the bus inactivity limit does not apply. - 0x0000 indicates the device should use the shortest value implemented. And - if the link is not a loop (no link tenancy) the value should be rounded to zero. Seems rounding to 0xFFFF would be more consistent.

The SPC-2 document specifies 0 as indicating no bus inactivity limit. There is no definition for the 0xFFFF value. The text of 9.1.1.3 is modified to correspond to the SPC-2 definitions. Note that the present proposed values came out of a white paper provided early in the FCP-2 work, and there may be legacy implementations that followed this definition.

[done](#)

3.14 LSI 015 * Disconnect Time Limit clarification (Technical)

page 48, 9.1.1.4 Disconnect time limit This is inconsistent with previous limit value (bus inactivity). Above a value of 0xFFFF meant does not apply, yet here we have 0x0000 meaning does not apply . Is there some justification for the inconsistent definition?

This is consistent with the SPC-2 definition and remains unchanged.

[no change for this comment](#)

3.15 LSI 016 Burst size during write clarification

page 48, 9.1.1.6 Maximum burst size field The text indicates this is the maximum a device shall transfer . I believe this also applies to write data where the device is requesting, not transferring data. I d suggest adding or request in an FCP_XFER_RDY IU to the end of the first sentence.

The suggested change is accepted.

[done](#)

3.16 LSI 017 * First Burst Size clarification (Technical)

page 49, 9.1.1.10 First burst size field The third paragraph of this section specifies the value shall be implemented by all FCP-2 device. This field is meaningless if the device doesn t support Write XFER_RDY Disable and never sets the bit in PRLI accept. I d suggest changing to value shall be implemented by all FCP-2 device which support Write XFER_RDY Disable.

The suggested change is accepted.

[done](#)

3.17 LSI 018 * Choice of RR_TOV default (Technical)**

page 53, Table 28 Why is RR_TOV so huge? 5 minutes may be the correct value for tapes (where this table was copied from) but I don t think it is appropriate for all devices (including disk and disk arrays). I d suggest a value in the 4 to 10 second range with a note deferring to any applicable profile for the device.

PLDA indicates that the default value when no recovery operations are allowed is 2 seconds. Otherwise, FCP-2 would require the default value to be $\geq 3 \times \text{REC_TOV}$. Note that 2.3.5 requires REC_TOV to be implied from RR_TOV. The table is modified to clarify these relationships, not requiring 300 seconds.

done

The editor will review the text to be sure that the case where no recovery is running is defined correctly.

3.18 LSI 019 * Clarify FLOGI (Technical)

page 53, Table 28 Note 1 wording implies this is FC_TAPE, not FCP-2. Note 3 correctly says devices shall use the value obtained in the FLOGI ACC after Fabric Login. However, the last sentence seems to contradict this. I believe the last sentence of note 3 should be deleted.

The suggested change is accepted.

done

3.19 LSI 020 * Class 1 and 4 treatment (Technical)

page 56, 11.1.1, 11.1.2 I don't what to define it, but can a standard ignore class's 1 and 4 like the profile did?

The consensus is that Class 1 and Class 4 are acknowledged classes and can be treated in large measure the same as Class 2. Within an exchange, FC-FS supposedly prohibits mixing of acknowledged and unacknowledged classes. The glossary of FC-FS may already have definitions for the acknowledged and unacknowledged classes. Note that REC and the exchange being recovered may be in different classes. The text of FCP-2 Revision 3 will be modified to define these relationships. Note that the word "delivery" in "class x delivery service" should be removed.

done, 3.1.1, 11

3.20 LSI 021 * Missing ACK after FCP_CONF (Technical)

page 56-57, 11.1.2 Need to cover the case where no ACK is received after FCP_CONF.

The suggested change is accepted. This is actually covered by 11.1.1, but is mislabeled as applying only to class 3 service. In fact, these detection mechanisms apply to all classes of service.

done 11.1.1, 11.1.2

3.21 LSI 022 ** ACK(abort) to target (Technical)

page 57, 11.1.2 The second to last paragraph of this section describes what the initiator should do if it receives an ACK with the Abort bits set. What should a target do if it receives the same?

This actually describes what the sequence initiator shall do. The sequence initiator may be a target or an initiator.

no change

3.22 LSI 023 Editorial

page 57, 11.2.1 Error Detection Confusing title for section. Previous 3 section are also flavors of error detection.

The suggested change is accepted. Titles for 11.1.1, 11.1.2, and 11.2.1 changed.

[done](#)

3.23 LSI 024 Editorial

page 57, 11.2.1 The wording flow from paragraph 1 to 2 is rough. Paragraph 2 needs to be reworded.

The suggested change is accepted.

[done](#)

3.24 LSI 025 Read XFER_RDY obsolete

page 57, 11.2.3 FCP_XFER_RDY Recovery Title needs to change to indicate Write XFER_RDY. Also, do we need a section concerning Read XFER_RDY?

In the meeting of September 14, 1999, there was unanimous consent that FCP_XFER_RDY for read operations should be made obsolete. PRLI parameters are changed to indicate that the control bit shall be set to always prohibit FCP_XFER_RDY on reads.

[done in 4.2, 5.4 \(Table 8\), 6.2.6 \(Table 10\), 6.2.6.12, 6.2.7 \(Table 11\), 8.2, 8.2.1, C.1.1, C.1.4, C.1.6, C.3,](#)

3.25 LSI 026 * Permission to do recovery (Technical)

page 57-58, 11.2.3 FCP_XFER_RDY Recovery This and following sections which indicate sending SRR need to qualify it with whether or not the Retry bit was on in PRLI.

The suggested change is accepted. A new section 11.1 is defined to explain how to do exchange level error recovery. Significant changes were required throughout section 11 to clarify this relationship.

[done, major changes to section 11. please review.](#)

3.26 LSI 027 * REC data count (Technical)**

page 58, 11.2.4 FCP_RSP Recovery Under 2) at the beginning of the section: The wording implies that the REC data count on a write is how much was transferred on the sequence. I thought the REC data count accumulated for the entire I/O. (I m thinking of the case where the target asks for the data in chunks with multiple XFER_RDYs.) We need to clarify REC and (if appropriate) fix the words here.

The intent is that the REC data count apply to the whole I/O. The change is accepted.

[done](#)

3.27 LSI 028 ** Recovery of 0 Bytes Xferred case (Technical)

page 58, 11.2.4 FCP_RSP Recovery Under 2) at the beginning of the section: As discussed last meeting there is no way to tell the difference between a lost XFER_RDY on a write command and a lost FCP_RSP in the case where the target was going to give a Check Condition and request CONF. In both cases 0 data bytes have been transferred and the host holds SI (according the target). We need to fix REC.

The target performs the correct recovery based on its known state and capabilities. Additional explanatory text was provided to make clear that the target was allowed to interpret the SRR suggested retry and correct it if necessary.

[done](#)

3.28 LSI 029 * Hold exchange information (Technical)**

page 58, 11.2.4 FCP_RSP Recovery Under 2) at the end of the section: The target is suppose to hold exchange info until n+1 commands have been received. This only works if Precise Delivery is enabled. Also, there is no way for the value of n to be known to both the init and target (in today s SCSI), yet this is vital information for this algorithm to work.

This created a number of subsidiary discussions in the September 14 meeting.

After significant discussion, the group decided that FCP_CONF could be requested at any time by the target. I assume that this applied only to FCP_RSP for SCSI commands, and not for task management functions. Dal Allan provided a model for the use of FCP_CONF which has not yet been reviewed by the committee, but will be included in revision 3 for review.

The group also decided that FCP_CONF should be a requirement for all queued commands while the target and initiator were operating with recovery behavior operating.

[done \(see 3.4\)](#)

The group further decided that SRR requesting status could be rejected if no status were available, just like it could be rejected for data recovery cases.

[done \(11.3.4 and 7.1\)](#)

Chapter 11 needs to further clarify what behaviors are optional when error recovery is in effect.

[done](#)

The value of N+1 is not the right number of exchanges to hold. The value will be reviewed.

[done](#)

3.29 LSI 030 ** RO during recovery (Technical)

page 58, 11.2.5 FCP_DATA Recovery Write The last sentence (above the note) says the SRR contained an RO. However that is not how SRR is currently defined. The RO field is only valid if requesting data, in this case

we are requesting a data descriptor. However, because Writes have a handshake between every data IU, I don't think it is a problem to not have the RO specified. The implication is retransfer the last write data IU. Need to change the words in this paragraph.

SRR is presently defined incorrectly. RO is presently defined only for solicited data retry. It is now clarified that it applies for an XFER_RDY missing after a solicited data transfer as well. It will also be clarified that RO is that of the lowest unsuccessfully transmitted data.

done (see 7.1 and 11.3.5)

Note that the FCP_XFER_RDY may have to be sent for a value lower than the SRR if the device must begin the transfer from some other boundary.

done (see 11.3.5) (please review)

3.30 LSI 031 * FCP_RSP retransmission (Technical)

page 59, 11.2.6 FCP_DATA Recovery Read Wording in second paragraph assumes the target has already transferred FCP_RSP once. This may not be true. Current words: The Target shall retransmit the requested data in a new Sequence, and then retransmit the FCP_RSP. Suggested new words: The Target shall retransmit the requested data in a new Sequence, and then complete the I/O as normal, including transmitting or retransmitting the FCP_RSP.

The suggested change is accepted.

done

Note that the target may need to send data beginning at a lower value than the RO specified in the SRR.

done (see 11.3.6) (please review)

3.31 LSI 032 ** Recovery RO (Technical)

page 59, 11.2.6 FCP_DATA Recovery Read Wording in second paragraph implies target may modify the RO. The current SRR definition requires the target to start at the specified RO.

In a number of discussions, we have concluded that a target may have to retry the entire data transfer from the beginning if it cannot begin at the specified location. This is at least one case where the RO must be modified by the target. The current SRR definition will be corrected to allow such behavior.

done

3.32 LSI 033 * OX_ID reuse and RX_ID (Technical)**

page 59, 11.2.7 FCP_CONF Recovery This algorithm assumes the target is assigning an RX_ID and the initiator is checking when it receives an REC. If not, we have a problem because the initiator may have already reused the OX_ID because it thought the I/O was complete. In other words, the REC ACC may be ambiguous to the target if the assigned RX_ID was not checked by the initiator.

After considerable discussion, it was concluded that an RX_ID needed to be provided and checked if FCP_CONF is used. It was further concluded that the use of an RX_ID of 'FFFFFF'X be prohibited in FCP-2 when confirmed completion is active.

done

Note that the OX_ID/RX_ID cannot be duplicated in that case and the last sentence must be modified to reflect that.

done (please review)

3.33 LSI 034 * OX_ID and RX_ID in REC (Technical)**

page 74, B.3.1 Read Exchange Concise (REC) The paragraph at the top of the page tells initiators to check the OX_ID and the targets to check the RX_ID and both to ignore the other half of the X_ID. I believe both sides need to check the entire X_ID. (see LSI 033 for an example)

Se3.32.

done

Just curiosity, but why is the S_ID required in the REC Request payload? Is any device other than the D_ID or S_ID of the REC going to be the originator?

Please review

4.0 Consideration of HP comments for inclusion in rev 3

The following comments were received from HP after the last meetings. Proposed changes to FCP-2, to be included in revision 03, are offered for each comment.

4.1 HP 1: New glossary terms

Clause 3.1 definitions: Add the terms "Recovery Abort" defined in clause 8.1.2.2, page 37 and "Ambiguous Exchange" defined in clause 8.1.1.4, page 35. Both of these terms are used in the text before they are formally defined.

The proposed changes are accepted in principal, although other clarifying methods may be used. The glossary is probably not the right place for the definition of such a complex function.

The word "ambiguous state" is defined one paragraph from its first use. A forward reference is supplied.

done

The word "recovery abort" is summarized in the glossary and a reference to the detailed definition is provided there.

done

4.2 HP 2: Reference for FQXID

Clause 3.1.16 fully qualified exchange identifier: include reference to clause 5.1 page 14 where it is fully defined.

The proposed change is included in revision 3 of FCP-2.

done

4.3 HP 3: References updated for documents

Clause 4.1, three references to ANSI X3.270 (SAM). Should these references be updated to SAM and SAM-2 or just SAM-2.

The references should be updated. The proposed change is included in FCP-2 revision 3.

done

4.4 HP 4: editorial

.Clause 4.6, page 10, last sentence: "table 2", should be "Table 2"

The proposed change is included in FCP-2 revision 3.

done

4.5 HP 5: Clarify implicit PRLI/PRLO is beyond scope

Clause 4.9, page 13, last sentence: "Implicit PRLI/PRLO parameters may be defined for FCP_Ports." Since this technique is not defined here should a sentence be added to state that this is beyond the scope of this document?

The proposed change is included in FCP-2 revision 3.

done

4.6 HP 6: editorial

Clause 6.2.6.5 and 6.2.6.6, page 24: Reader might assume that annex

A is part of ANSI X3.230. Clarify text to state "annex A in this document".

The actual reference is ANSI X3.297, FC-PH-2.

done

4.7 HP 7: editorial

Clause 6.2.6.9, page 24, third line: Cut and paste error? "FCP_DATA IUs that are move data".

The proposed change is included in revision 3 of FCP-2.

done

4.8 HP 8: editorial

Clause 6.2.6.10, page 24, last line: missing period at end of paragraph.

The proposed change is included in revision 3 of FCP-2.

done

4.9 HP 9: Correct directional reference

Clause 11.1.2, page 57 in "The Exchange responder (SCSI Target).... list line "4) a Sequence error is detected in an inbound Sequence" should be an outbound sequence. (The target receives only outbound sequences.)

The proposed change is included in revision 3 of FCP-2.

done

4.10 HP 10: Target retry of FCP_CONF

Clause 11.2.7 FCP_CONF Recovery: This clause should clearly state that in the case of a lost FCP_CONF, that the Class 3 target is required to initiate recovery action. This is different than in the cases described in clause 11.2.5 FCP_DATA Recovery - Write and 11.2.9 Additional error detection by SCSI Target where the text explicitly states that the class 3 target does not initiate error recovery.

The proposed change is included in revision 3 of FCP-2.

[done](#)

4.11 HP 11: editorial

Annex A.4, page 69: Service response description has a formatting problem. At least it is missing a square bracket.

The proposed change is included in revision 3 of FCP-2. The square bracket is actually not supposed to be there.

[done](#)

4.12 HP 12: Correct references and table A.1

Annex C, page 77: The tables are labeled with "A" instead of "C". Table A.1 has the arrow pointing in the wrong direction on "I3 FCP_DATA".

The proposed change is included in revision 3 of FCP-2.

[I have verified the arrow direction, and it appears to be correct.](#)

[done](#)

4.13 HP 13: Correct references

Annex D, page 87: The figures are labeled with a "C" instead of a "D".

The proposed change is included in revision 3 of FCP-2.

[done](#)

4.14 HP 14: Correct figure C.9

Annex D, page 95, "Figure C.9 - Lost Read Data, Last Frame of Sequence"

The Class 3 Error detection drawing has the REC and ACC arrows in the wrong direction.

[After review of this figure, I believe the error directions are correct. The REC is the result of a sequence error detection and is generated by the initiator.](#)

5.0 Review actions required

5.1 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

In the July meeting, there was considerable discussion about this question. Dave Peterson feels that most of the work required to make out-of-order operation behave correctly is already included. Dal Allan and Carl Zeitler believe that out-of-order operation should be allowed if at all possible. Please review section 11 carefully for discrepancies that may cause failures during recovery of out-of-order transfers. If there are none, we can remove the “in-order-only” restriction.

5.2 Combined command/data, data/response obsolete

This was made obsolete in revision 02. Some parts were not corrected. The following portions of the document required additional work.

[C.1.7, removed](#)

[C.1.8, removed](#)

6.0 Other stuff to worry about

6.1 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 6.2, which proposes that process associators be made obsolete in FCP-2.

not done in revision 03

6.2 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.

The resolution depends on review work being done by at least one company on the possible uses of Process Associators.

not done in revision 03

6.3 Incorrect use of Recovery Abort

Section 11.4.1 (now 11.5.1) defines a number of cases where the recovery abort protocol is supposed to be executed. Many of these definitions are incorrect. In particular, ABTS should not be invoked following CHECK CONDITION status for resets or microcode changes. These are normal SCSI behaviors. Many of them conflict with section 8.1.4, which requires ABTS-LS only for exchanges whose state is ambiguous. I propose that we do the following:

- 1) a-1 should apply only to ambiguous exchanges.
- 2) a-2, both sections should be deleted.
- 3) b-1 and b-2 should apply only to ambiguous exchanges.
- 4) b-3 should be deleted.

to be reviewed

6.4 Specification of formats for ELS

Section 11.6 (now 11.7) is a hold-over from profile days. Should we move this to an informative annex? It specifies the details of the contents of the ELSs used in the recovery procedures, probably in a redundant manner with respect to FC-FS and other documents.

to be reviewed

6.5 Multi-initiator

Section 11.8 (now 11.9) contains some multi-initiator behavior definitions. This should be moved to an informative annex. It specifies behaviors that may conflict with SAM-2 and SPC-2.

to be reviewed