

Accredited Standards Committee

NCITS, National Committee for Information Technology Standards

Doc: T10/97-266r1

Date: 7/14/98

Project: FPC-2

Ref Doc:

Reply to: Bob Snively

To: T10 Membership

From: Bob Snively

Subject: FCP-2 Initial Changes

There have finally been enough minor changes and corrections proposed for FCP that it is time to collect them together and begin the editing work for FCP-2. This document collects references to the known FCP-2 modifications and is intended to become an exhaustive list of the initial changes to be made in the document, including major editorial changes and technical changes.

1.0 Editorial Changes

1.1 Obtain ANSI edits

The ANSI editor has provided numerous updates to references and small editorial modifications which must be obtained for inclusion as part of FCP-2.

1.2 Provide complete mapping of service interfaces.

Provide a clause similar to clause 11.8 in SPI-2 (T10/1142D revision 15) to identify the mapping between the services defined by SAM (or SAM-2 if applicable) and FCP. These definitions may replace one or more paragraphs scattered through the document.

1.3 Remove Annex A

The contents of Annex A are now documented by one of the FC standards. After appropriate review for completeness, the annex will be removed and the proper FC documents referenced.

1.4 Correct bit definition

The last sentence on page 41 of revision 12 should refer to bit 13, not bit 14. (Chan, July 24, 95)

1.5 Correct PRLI Accept Response code

Page 19 specifies that the PRLI Accept Response code of 1000 is Invalid Service Parameters for page. However, the FC documents have a different value. This should be corrected. (Snively/Binford, Sept 18, 95.)

1.6 Clarify FCP_RSP formats for task management

There is some lack of clarity about how task management responses are implemented. The text should be improved. (Frazier, 7 Feb 96)

1.7 Clarify FCP ABTS when no exchange exists yet

Charles Monia (5 Aug, 96) notes that the ABTS responses do not correctly address the case where the ABTS may precede the arrival of the exchange to be aborted. He provides a recommended clarification which will be reviewed.

1.8 Definition of Data Overlay

Ed Gardner (30 Aug, 96) notes that data overlay is not defined in the present FCP document. A correct definition will be provided. A series of mails in September of 1997 also address this issue.

1.9 Definitions that should be included

George Penokie (19 Sept, 1997) has requested the following clarifications:

- 1) Many of the acronyms are not defined or defined in places where the definition is difficult to find. Examples include: OX_ID, RX_ID, OOA, ROA, PLOGI and perhaps others.
- 2) "Process association" should be defined.

1.10 Removal of levels of indirection

George Penokie (19 Sept, 1997) suggests that the document would be clearer if tables describing IUs were integrated, instead of providing a hierarchy that must be interpreted. He mentions the FCP_CMND field and the FCP_CNTL field as examples. Where appropriate his suggestion will be followed.

1.11 Clarify usage of XFER_RDY during read

Gen-Hwa Chiang (29 Oct 97) asks about the proper operation of multiple sequence reads when XFER_RDY is not used. I will examine the document to see if this is clearly stated and clarify this if not.

2.0 Technical Changes

2.1 95-348r1, FCP usage of Disconnect Reconnect Page

This proposal has a number of technical additions proposed for FCP. I will review the proposals, the minutes of the appropriate meetings, and the actual implementations and provide proposed text. It may be that, after later review, some of the proposals should be left out of FCP. Several of these functions are also required by PLDA.

Items considered by this document include:

1) Use of disconnect reconnect page parameters

- Buffer full and empty ratios
- Bus inactivity limit field
- Disconnect time limit
- Connect time limit
- Maximum burst size field
- Enable Modify Data Pointers
- Access fairness control bits
- Disconnect Immediate bit
- Data transfer disconnect control field
- First burst size field

2.2 96-195r4, FCP control page parameters

This proposal describes a number of technical additions proposed for FCP. I will review the proposals, the minutes of the appropriate meetings, and the actual implementations and provide proposed text. It may be that, after later review, some of the proposals should be left out of FCP. These functions are also mentioned by PLDA. Mail from John Nutter (2 May, 1997) further places clarification requirements on the text to be inserted.

Items considered by this document include:

- ### 1) Fibre Channel Control page, including
- Disable Target Originated Loop Initialization
 - Disable Target Initiated Port Enable
 - Allow Login Without Loop Initialization
 - Disable Soft Address
 - Disable Loop Master
 - Disable Discovery

2.3 Add ABORT LOGICAL UNIT task management function

This has now been defined in SAM-2 and will be added to FCP-2.

3.0 Technical changes to be discussed and approved

3.1 Flag bit usage

A Spaldin (14 June 96) has requested clarification of the flag bit in FCP linked commands. At present, there is no definition in SCSI as to the mandatory nature of the flag bit. Since no host adapter actually uses either linking or the flag bit, I would like to start phasing it out of SCSI, starting with FCP.

3.2 Serial link usage of RESERVE/RELEASE(6)

The present RESERVE(6)/RELEASE(6) only has 3 bits of identifier for third-party functions. I would prefer to make those commands obsolete for FCP devices and allow the use of RESERVE(10)/RELEASE(10) and PERSISTENT RESERVE IN/OUT instead.

3.3 RR_TOV timer

Jim Coomes references a PLDA annex (Annex D) that will contain 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.

3.4 Clearing of mode pages

PLDA specifies a number of actions in table 16 that are forced by Fibre Channel operations with respect to initialization states, mode pages, tasks, and task sets. These will be reviewed to see if any need to be included in FCP-2.

3.5 Definition of mode page parameters for FCP

PLDA specifies in table 18 a selection of interpretations for mode sense/select information. These will be reviewed to see if any need to be included in FCP-2.

3.6 Use of mode page settings to control initialization

PLDA specifies in section E.5.1 the possibility of controlling initialization using mode pages. This will be identified and reviewed to see if any need to be included in FCP-2.

3.7 TERMINATE TASK removal

At present, the TERMINATE TASK function is still included in FCP. With its probable promotion to obsolete in SPI-2, it will also be made obsolete or removed in FCP-2.

3.8 Clarification of FCP_CDB content

George Penokie (19 Sept, 1997) indicates that the statement in section 7.1.3, "The FCP_CDB is not valid and is ignored if any task management flag is set to 1" is not correct. He believes that Clear ACA should be allowed along with a valid CDB. While this may be true in parallel SCSI, it is not a requirement of SAM and there does not appear to be a need for it in FCP. I propose that no change be made.

3.9 Command Reference Number

This has been proposed to provide quick identification of missing commands, especially important in a queued tape environment where ordering may be demanded by the task attribute. This has been approached in FC-TAPE with the following wording:

The FCP Command Reference Number is used to ensure proper ordering of Exchange's (SCSI commands). The use of CRN is enabled by setting the Command Reference Number Enabled bit to one in the PRLI request. See Annex H for a description of the CRN bit.

The CRN is contained in Byte 0 of the FCP_CNTL field in the FCP_CMND IU and shall be continuously increasing. The first CRN is equal to one for the first Exchange between the Initiator and Target based on an I_T_L nexus. A value of zero in the CRN field is only valid for an existing Exchange (i.e. the CRN counter has wrapped). A PRLI and Target Reset Task Management Function shall reset the CRN to zero.

The Initiator shall not reuse a CRN until delivery has been confirmed via a FCP_XFER_RDY, FCP_DATA, FCP_RSP IU or a REC or ACK.

The Target shall not execute out-of-order SCSI commands and shall hold the SCSI command until prior CRN(s) have been received.

There is still a great deal of debate about this issue and that debate must lead to a resolution before this can be properly included in FCP-2. When resolved, this will go into FCP-2.

3.10 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 should be defined in FCP-2 and referenced in FC-TAPE. At present, the FC-TAPE document is ahead of the FCP-2 document, so I am hoping that the text that will go into FC-TAPE will be formatted as an informative annex so that it can be standardized in FCP-2.

3.11 FCP_CONF

The present FC-TAPE document proposes the implementation of an FCP_CONF (Confirm) IU that would be requested by an FCP_RSP and returned by the initiator to the target to inform the target that its response has been received. This has the possible attributes of synchronizing the state of the initiator and the target/LUN more accurately and allowing for the recovery of status information that was not transmitted correctly. FC-TAPE has the following words about this at present.

The FCP_CONF IU is used to confirm receipt of a FCP_RSP. The use of FCP_CONF is enabled by setting the FCP_CONF Enabled bit to one in the PRLI request. A Target device conforming to this profile is required to support the reception of a FCP_CONF IU if the Target device supports tagged command queuing using Class 3 delivery service.

The FCP_CONF is sent by the Initiator after a FCP_RSP has been received with the FCP_CONF_REQ bit set in the FCP_STATUS field. The Initiator shall release Exchange information such as the Exchange Status Block (ESB) after the FCP_CONF is sent. The Target shall retain Exchange information and associated data until a FCP_CONF is received. See Annex G for a description of the FCP_CONF_REQ bit and FCP_CONF usage.

There is still a great deal of debate about this issue and that debate must lead to a resolution before this can be properly included in FCP-2. When resolved this will go into FCP-2.

3.12 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.

3.13 Parameter associated with initiator

Charles Monia (Apr 24, 1996) among others has asked how initiators are identified for reservations. At present, the initiator ID and Process Associator (if any) are the proper values. If those are re-assigned by LIP actions, then reconfiguration is necessary. Persistent reservation and software conventions should adequately manage this case, and FCP can't do much about it anyway.

An informal proposal has been circulated suggesting that a properly logged in Target has all the necessary information to couple a reservation to the initiator WWN. If it is specified that this information is used, then during post initialization target/initiator verification using FAN or PDISC/ADISC, a persistent reservation can actually be reconnected to the proper initiator independent of its AL_PA. If this proposal is accepted it will be included in FCP-2 as mandatory for those targets that implement persistent reservation.

4.0 Other suggestions, not planned for inclusion

4.1 First Burst Parameter

In e-mail dated Jan 3, 1996, Giles Frazier suggested that the first-burst parameter of the Disconnect-reconnect mode page be used to allow a different length limit for the first write data to be transmitted when the XFER_RDY function is not used. At present this function is not defined by FCP or used by any devices. I propose that it be excluded from FCP-2.

4.2 Concept of confirmation

The concept of using FCP IUs as confirmation, thus enabling the early re-use of SEQ_IDs within the same exchange is described as an implementation option (but not a requirement) in PLDA. I see that as an implementation alternative tutorial which should not be included in the FCP-2 document.

Sincerely,

Robert Snively

Phone:(415) 786-6694

Email: bob.snively@sun.com