Doc: T10/97-266r2
Date: 7/26/98
Project: FPC-2
Ref Doc:
Reply to: Bob Snively

To: T10 Membership

From: Bob Snively, Sun Microsystems

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. Revision 1 reflects the results of the FC-TAPE meeting of June 9, 1998. Revision 2 reflects the additional inputs from the FC-TAPE meeting of July 14, 1998 and the SCSI Working Group meeting of July 15, 1998.

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 the FC-PLDA document and by the FC-TAPE document, Annex I. Mail from John Nutter (2 May, 1997) further places clarification requirements on the text to be inserted.

The latest revision I have available indicates that the mode page is expected to have the following format and contents, although this text has not yet been verified against all the possible document sources.

Table 1: Fibre Channel 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							ECRN
5	Reserved							
6	Reserved							
7	Reserved							

When Disable Target Originated Loop Initialization (DTOLI) bit is one, the Target does not generate the Initializing LIP following insertion into the loop. The Target will respond to an Initializing LIP when it is received. The Target shall generate the Loop Failure LIP if it detects loop failure at its input and the Initializing LIP when the loop failure is corrected. When DTOLI bit is zero, the Target generates the Initializing LIP after it enables a port into a loop.

When Disable Target Initiated Port Enable (DTIPE) bit is one, the Target waits for an Initiator to send the Loop Port Enable primitive before inserting itself into the loop. The Target uses the hard Address available in the SCA connector or device address jumpers to determine if primitives are addressed to it. A Loop Port Enable primitive with the broadcast address shall also cause the Target to insert itself into the loop. When DTIPE bit is zero, the Target enables its port into the loop without waiting for a Loop Port Enable primitive.

Wen Allow Login Without Loop Initialization (ALWLI) bit is one, the Target shall use the hard address available in the SCA connector or device address jumpers and accept logins without verifying the address with loop initialization. When ALWLI bit is zero, the Target is required to verify its address through the loop initialization process before a login is accepted.

When Disable Soft Address (DSA) bit is one, the Target does not select a soft address if there is a conflict for the hard address selection during loop initialization. In this case the Target enters the nonparticipating state. If the Target detects loop initialization while in the nonparticipating state, the Target will again attempt to get its hard address. When DSA bit is zero, the Target attempts to obtain a soft address during the loop initialization process.

When Disable Loop Master (DLM) bit is one, the Target does not become loop master. The Target only repeats LISM frames it receives. This allows the Initiator to be loop master during loop initialization. When DLM bit is zero, the Target may become loop master during the loop initialization process.

When Disable Discovery (DDIS) bit is one, the Target does not require receipt of Address or Port Discovery following loop initialization. The Target resumes processing of tasks on completion of loop initialization. When DDIS bit is zero, the Target must wait to receive an Address or Port Discovery before it resumes processing tasks for that Initiator.

When the Prevent Loop Port Bypass (PLPB) bit is one, the Target shall ignore any Loop Port Bypass (LPB) and Loop Port Enable (LPE) primitive sequences. The loop port shall remain enabled. When PLPB bit is zero, the Target allows the Loop Port Bypass (LPB) and Loop Port Enable (PBE) primitive sequences to control the port bypass circuit.

It shall be illegal to set DTIPE to one and PLPB to one.

When an illegal bit combination is sent by the application client the device server shall return CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN THE PARAMETER LIST.

When Enable Command Reference Number (ECRN) is one, the Initiator shall use a continuously increasing Command Reference Number for each Exchange with the Target. When the ECRN bit is zero, the Initiator shall not use a continuously increasing Command Reference Number.

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 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.

3.3 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. It has been recommended by the working group that this information be contained in a table of “event equalities” between SCSI and FCP.

3.4 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.

3.5 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.

3.6 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.

3.7 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.

3.8 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, updated as of the July 14 meeting. The intent of the FC-TAPE document is that annex D will be supplanted by FCP-2.

8.2.1.1 FCP Command Reference Number (CRN)

The FCP Command Reference Number shall be used to ensure proper ordering of Exchange's (SCSI commands). The Target shall indicate support of CRN by setting the Command Reference Number Support (CRNS) bit to one in the PRLI request and enable CRN usage on a per-Lun basis by setting the Enable Command Reference Number (ECRN) bit to one in the FC Mode Page (0x19) for the Lun. See Annex D for a further description of CRN implementation and usage.

Annex D (informative)

Command Reference Number (CRN) Implementation and Usage

This annex describes the implementation and usage of Command Reference Number (CRN).

The CRN is contained in Byte 0 of the FCP_CNTL field in the FCP_CMND IU and is based on an I_T_L nexus. Support for CRN is negotiated based on a I_T nexus via PRLI. The use of CRN is enabled based on a I_T_L nexus via FC Mode Page (0x19).

D.1 Command Reference Number Process Login (PRLI) Parameter

D.1.1 Word 3, Bit 8: Command Reference Number Support (CRNS)

When this bit is set to 1, the process defined by the page is indicating that it is capable of supporting CRN. When this bit is set to 0, the process defined by the page is indicating that it is not capable of supporting CRN. If the responder does not support CRN it shall set the bit to 0 in the PRLI accept page.

D.2 FC Mode Page (0x19) Parameter

D.2.1 Byte 4, Bit 0: Enable Command Reference Number (ECRN)

When the Enable Command Reference Number (ECRN) bit is one, the Initiator shall use a continuously increasing Command Reference Number for each Exchange with the LUN (except Task Management Functions). When the ECRN bit is zero, the Initiator shall not use a Command Reference Number for each Exchange with the LUN.

D.3 Guidelines

1) The CRN shall be equal to 1 for the first Exchange between the Initiator and Target and shall be continuously increasing.

2) The CRN shall wrap from 255 to 1 (i.e. a value of 0 in the CRN field is not valid for an Exchange using CRN).

3) A PRLI, Target Reset Task Management Function, and LUN Reset Task Management Function shall reset the CRN to 1.

4) 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.

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

6) Task Management Functions shall not use a CRN and shall set the CRN field to 0.

7) Commands with a CRN of zero shall not participate in the ordering and verification procedures specified in 5.

Any command, including such initialization commands as INQUIRY, TEST UNIT READY, and MODE SENSE/SELECT may always use a CRN of zero if the state of the ECRN bit is not known or if execution ordering is not required for that command.

7.1 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.

7.2 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. This has been approached in FC-TAPE with the following wording, updated as of the July 14 meeting. The intent of the FC-TAPE document is that annex C will be supplanted by FCP-2.

8.2.5 FCP_CONF (T12)

The FCP_CONF IU is used to confirm receipt of a FCP_RSP. Support for FCP_CONF is negotiated using FCP_CONF Support bit to one in the PRLI request. A Target device conforming to this profile shall request a FCP_CONF IU for each FCP_CMND containing a Task Attribute of SIMPLE_Q, HEAD_OF_Q, or ORDERED_Q (i.e. tagged task).

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 C for a description of the FCP_CONF_REQ bit and FCP_CONF usage.

Annex C

(informative)

FCP_CONF Implementation and Usage

This annex describes and diagrams the use of FCP_CONF.

SCSI Initiator and Target devices conforming to this profile are required to support FCP_CONF IUs.

C.1 FCP_CONF Information Unit

[Table 27 provides the definition for IU T12, to be placed in the proper FCP table.]

C.2 FCP_CONF Process Login (PRLI) Parameter

C.2.1 Word 3, Bit 7: FCP_CONF Support

When this bit is set to 1, the process defined by the page is indicating that it has the capability of supporting FCP_CONF IUs. When this bit is set to 0, the process defined by the page is indicating that it is not capable of supporting FCP_CONF IUs. If the responder does not support FCP_CONF IUs it shall set the bit to 0 in the PRLI accept page.

C.3 FCP_CONF Request Returned in FCP_RSP

A FCP_CONF request may be sent by the Target in the FCP_RSP by setting the FCP_CONF_REQ bit in the FCP_STATUS field. FCP_CONF_REQ indicates the Target has requested a FCP_CONF for the Exchange.

C.4 FCP_CONF Guidelines

1) The Initiator shall send a FCP_CONF IU after receiving a FCP_RSP from the Target with the FCP_CONF_REQ bit set in the FCP_STATUS field.

2) The Initiator shall release Exchange information after the FCP_CONF is sent or when a response to an Exchange abort is received.

3) The Target shall retain Exchange information until a FCP_CONF is received or the Exchange is aborted.

4) ABTS shall be sent by the Target if no FCP_CONF IU is received by the Target within FCP_TOV after a FCP_RSP has been sent.

5) The use of FCP_CONF is required for Tagged Tasks (i.e. command queuing). [Editor's note: This may not be a requirement and will be discussed in subsequent meetings.]

[Figure 21 Shows a ladder diagram providing an example of the use of FCP_CONF. It is described in the adjacent text, below.]

A FCP_RSP requesting a FCP_CONF is sent by the Target. The FCP_CONF is sent by the Initiator and received by the Target within FCP_TOV.

[Figure 22 Shows a ladder diagram providing an example of the behavior of FCP_CONF recovery if the FCP_CONF is not received. The following text describes the actions.]

If a FCP_CONF is sent by the Initiator and not received at the Target the timer will expire and the Target sends an ABTS. The Initiator does not know of the Exchange and returns a BA_RJT with an "Invalid OX_ID-RX_ID combination" reason code explanation. The Target knows the FCP_CONF was sent and releases the Exchange information.

If another FCP_CMND is received with the same OX_ID as an Exchange waiting for a FCP_CONF the Target may release the Exchange information.

There is still considerable work that must completed for FCP_CONF to become a useful and usable addition to the standard.

5.1 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.

5.2 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

reassigned 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.

A proposal (T10/98-206r1) 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.

6.0 Items not planned for inclusion

6.1 Concept of confirmation

The concept of using FCP IUs as confirmation, thus enabling the early reuse 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