Doc: T10/97-266r3
Date: 10/30/98
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 approved FCP-2 modifications. It is an exhaustive list of the initial changes 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. Revision 3 reflects the actual changes made in the document during the initial update from FCP to FCP-2.

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.

Done. Rev 00 is based on ANSI edits, updated to reflect present references.

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.

(not yet completed)

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.

Done. FC-PH-2 section 21.10, 21.11, and 21.12 were verified to contain all technical information contained in annex A.

1.4 Remove Annex C

The LU addressing models of Annex C are now defined as mandatory SCSI behavior in SAM-2.

Annex C is removed.

1.5 Remove Annex D

The bibliography contained in Annex D is now allowed into the Normative References clause of the text, clause 2. As a result, the contents of Annex D, where still valid, are transferred to clause 2.

1.6 Correct bit definition

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

Done. This is in Annex A and has been removed.

1.7 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.

Done: The PRLI and PRLO Accept Response code of 1000 should be Invalid Service Parameters. Annex B is created to contain this and similar corrections to be applied to other documents.

1.8 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)

(not yet completed)

1.9 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.

(Not yet completed)

1.10 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.

(Not yet completed)

1.11 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.

(Not yet completed)

1.12 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.

(Not yet completed)

1.13 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.

(Not yet completed)

1.14 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.

(Not yet completed)

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 have reviewed the proposal and made significant modifications in the wording. The proposal goes into section 8.1.2 of the document.

8.1.2 Disconnect-Reconnect mode page

The disconnect-reconnect page (see Table 1) provides the application client the means to attempt to optimize the performance of the service delivery subsystem. The name for this mode page (disconnect-reconnect) comes from the SCSI-2 parallel bus. An SCSI-3 FCP-2 device uses only the parameters described in this clause. The application client and initiator communicate to determine what values are most appropriate for a device server. The device server communicates the parameter values in this mode page to the Target Role Agent, normally the Fibre Channel interface circuitry. This communication is internal to the target device and is outside the scope of SCSI-3. If a parameter that is not appropriate for the an FCP-2 SCSI-3 device is non-zero, the device server shall return CHECK CONDITION status.

The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

Table 1 - Disconnect-reconnect page (02h)

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Resvd	Page Code (02h)					
1	Page Length (0Eh)							
2	Buffer full ratio							
3	Buffer empty ratio							
4	(MSB)	Bus inactivity limit						
5								(LSB)
6	(MSB)	Disconnect time limit						
7								(LSB)
	(MSB)	Connect time limit						
								(LSB)
	(MSB)	Maximum burst size						
								(LSB)
	EMDP	FARd	FAWrt	FAStat	Reserved	Reserved		
	Reserved							
	(MSB)	First burst size						
								(LSB)

An interconnect tenancy is a period of time during which a target device owns or may access the interconnect. For example, on FC-AL loops or Fibre Channel Class 1 connections, a tenancy typically begins when a device successfully opens the connection and ends when the device releases the connection for use by other device pairs. Data and other information transfers take place during interconnect tenancies. Point-to-point or fabric-attached Class 2 or Class 3 links and many other configurations do not have a concept of link tenancy.

The buffer full ratio field indicates to the device server, during read operations, how full the buffer should be prior to requesting an interconnect tenancy. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in SPC-2. Devices attached to links that do not have the concept of link tenancy shall round the ratio to zero and transmit data in a vendor specific manner.

The buffer empty ratio field indicates to the device server, during write operations, how empty the buffer should be prior to requesting an interconnect tenancy. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in SPC-2. Devices attached to links that do not have the concept of link tenancy shall round the ratio to zero and request the transmission of data in a vendor specific manner.

The values contained in the buffer full ratio and buffer empty ratio are defined by SPC-2.

The bus inactivity limit field indicates the maximum duration of any interconnect tenancy during which no data is transferred, measured in transmission word increments. Target devices in configurations having the concept of interconnect tenancy shall end the tenancy if this limit would be exceeded. A value of zero indicates that the target device should use the shortest bus inactivity limit that it implements; the precise value is vendor unique. A value of all ones (0xFFFF) indicates that the bus inactivity limit does not apply. The device server may round this value to any value it prefers. Devices attached to links that do not have the concept of link tenancy shall round the bus inactivity limit to zero.

The disconnect time limit indicates the minimum delay between FCP interconnect tenancies measured in increments of 128 transmission words. Target devices in configurations having the concept of interconnect latency shall delay at least this time interval after each FCP interconnect tenancy before beginning arbitration. The device server may round this value to any value it prefers. A value of zero indicates that the disconnect time limit does not apply. Devices attached to links that do not have the concept of link tenancy shall round the disconnect time limit to zero.

The connect time limit field indicates the maximum duration of a single interconnect tenancy, measured in increments of 128 transmission words. If the connect time limit is exceeded the device server shall conclude the interconnect tenancy, within the restrictions placed on it by the applicable Fibre Channel configuration. The device server may round this value to any value it prefers. A value of zero indicates that there is no connect time limit. Devices attached to links that do not have the concept of link tenancy shall round the connect time limit to zero.

The maximum burst size field indicates the maximum amount of data that the device server shall transfer using a single Fibre Channel sequence. This parameter does not affect how much data is transferred in a single interconnect latency. This value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes, etc.). The device server may round this value down as defined in SPC-2. A value of zero indicates there is no limit on the amount of data transferred per data transfer operation. 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.

The enable modify data pointers (EMDP) bit indicates whether or not the target may reorder FCP_DATA IUs for a single SCSI command. If the EMDP bit is zero, the target shall generate continuously increasing DATA_RO values for each FCP_DATA sequence for a single SCSI command. If the EMDP bit is one, the target may transfer the FCP_DATA IUs for a single SCSI command in any order. An EMDP bit of zero prohibits data overlay, even if it is allowed by the state of the PRLI Data Overlay Allowed bit. This bit does not affect the order of frames within a sequence. The EMDP function is optional for all FCP-2 devices.

The FARd, FAWrt, and FAStat bits indicate whether a target in a loop configuration shall use the access fairness algorithm when beginning the interconnect tenancy. An FA bit of one indicates that the target shall use the access fairness algorithm. An FA bit of zero indicates that the target may choose to not use the access fairness algorithm. The FARd bit controls arbitration when the target wishes to send one or more FCP_DATA frames to an initiator. The FAWrt bit controls arbitration when the target wishes to send one or more FCP_XFER_RDY frames to an initiator. The FAStat bit controls arbitration when the target wishes to send one or more FCP_RSP frames to an initiator or FCP_CMND frames to another target. If the target intends to send multiple frame types, it may choose to not use the access fairness algorithm if any applicable FA bit is zero. Devices attached to links that do not have the concept of link tenancy shall ignore the FA bits. The FA bits are optional for all FCP-2 devices.

The disconnect immediate (Dimm) and the data transfer disconnect control (DTDC) fields defined in SPC-2 are not implemented and are reserved for FCP-2 devices.

2.2 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.

2.3 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.

The changes also provide the features casually mentioned in PLDA section E.5.

Done: The text is inserted in a new clause 8 that indicates controls that SCSI has over the FCP protocol. This text is significantly modified, since it refers principally to FC-AL connections. Many of the control bits are meaningless for SCSI devices connected in a fabric or on a point-to-point link.

The Fibre Channel Control mode page (see table 2) contains those parameters that select Fibre Channel operation options. The page shall be implemented by all FPC-2 devices. The implementation of any bit and its associated functions is optional. the page follows the MODE SENSE / MODE SELECT rules specified by SPC-2.

Table 2 - 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							

A Disable Target Originated Loop Initialization (DTOLI) bit of one indicates that a target attached by an FC-AL loop shall not generate an Initializing LIP following insertion into the loop. The target shall respond to an Initializing LIP when it is received. When DTOLI bit is zero, the target attached by an FC-AL loop shall generate the Initializing LIP after it enables a port into a loop. If the target is attached to an FC-AL loop and detects loop failure at its input, it shall follow the error initialization process defined by FC-AL-2. Targets attached to an N_Port or to an F_Port shall ignore this bit.

A Disable Target Initiated Port Enable (DTIPE) bit of one indicates that a target attached to an FC-AL loop shall wait for an initiator to send the Loop Port Enable primitive sequence before inserting itself into a loop. The target shall wait in a non-participating state with the Port Bypass circuit, if any, set to bypass the target. The target uses the hard address available in the SCA-2 SFF-8067 connector or in device address jumpers to determine whether LPE primitive sequences are addressed to it. A LPE primitive sequences addressed to the broadcast address shall also cause the target to insert itself into the loop. When the DTIPE bit is zero, the target shall enable itself onto the loop in a vendor specific manner. Targets attached to an N_Port or to an F_Port shall ignore this bit.

An Allow Login Without Loop Initialization (ALWLI) bit of one indicates that a target attached to an FC-AL loop shall use the hard address available in the SCA-2 SFF-8067 connector or in device address jumpers, enter the monitoring state in participating mode, and accept logins without using the loop initialization procedure. When the ALWLI bit is zero, the target shall perform the normal loop initialization procedure before entering the monitoring mode and accepting a LOGIN ELS. Targets attached to an N_Port or to an F_Port shall ignore this bit.

A Disable Soft Address (DSA) bit of one indicates that a target attached to an FC-AL loop shall not select a soft address if there is a conflict for the hard address selection during loop initialization. If there is a conflict, the target shall enter 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 the DSA bit is zero, the target follows the normal initialization procedure, including the possibility of obtaining a soft address during the loop initialization process. Targets attached to an N_Port or to an F_Port shall ignore this bit.

A Disable Loop Master (DLM) bit of one indicates that a target attached to an FC-AL loop shall not participate in loop master arbitration and shall not become loop master. The target shall only repeat LISM frames it receives. When the DLM bit is zero, the target may shall participate in loop master arbitration in the normal manner and may become loop master during the loop initialization process. Targets attached to an N_Port or to an F_Port shall ignore this bit.

A Disable Discovery (DDIS) bit of one indicates that a target attached to an FC-AL loop shall not require receipt of Address or Port Discovery (ADISC or PDISC ELSs) following loop initialization. The logical units shall resume processing tasks on completion of loop initialization. When the DDIS bit is zero, the target shall wait to complete target discovery as defined by FC-PLDA and FC-FLA before allowing processing of tasks to resume. Targets attached to an N_Port or to an F_Port shall ignore this bit.

A Prevent Loop Port Bypass (PLPB) bit of one indicates that a target attached to an FC-AL loop shall ignore any Loop Port Bypass (LPB) and Loop Port Enable (LPE) primitive sequences. The loop port shall always remain participating. 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 and participation on the loop as specified by FC-AL. Targets attached to an N_Port or to an F_Port shall ignore this bit.

The DTIPE and PLB bits shall not both be set to one at the same time. When an invalid 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.

An Enable Command Reference Number (ECRN) bit of one indicates that the target shall use the command reference number function defined by this standard. (See TBD). When the ECRN bit is zero, the target shall not use the command reference number function and shall ignore the contents of the CRN field. This bit is valid for all types of link connections. If the precise ordering function is not supported by the target, the ECRN bit shall remain zero even if set by a MODE SELECT command and the bit shall be masked as not changeable. (See SPC-2)

2.4 Add ABORT LOGICAL UNIT task management function

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

Done.

LOGICAL UNIT RESET, when set to 1, performs a reset to the logical unit as defined in SPC-2. LOGICAL UNIT RESET resets the logical unit and all dependent logical units as described below.

To execute a LOGICAL UNIT RESET the logical unit shall:

1) Abort all tasks in its task set(s);
2) Clear an auto contingent allegiance (NACA=1) or contingent allegiance (NACA=0) condition, if one is present;
3) Release all SCSI non-persistent device reservations;
4) Return the device's operating mode to the appropriate initial conditions, similar to those conditions that would be found following device power-on. The MODE SELECT parameters (see the SPC-2 standard) shall be restored to their last saved values if saved values have been established. MODE SELECT parameters for which no saved values have been established shall be returned to their default values;
5) Set a unit attention condition; and
6) Initiate a logical unit reset for all dependent logical units (see 4.11).

The FCP login state of the affected image pairs is not changed by the LOGICAL UNIT RESET.

The LOGICAL UNIT RESET is transmitted by the initiator (exchange originator) using a new exchange. Any open exchanges that are in an ambiguous state shall be terminated by whichever port detects the ambiguous state using a recovery abort. The ports may issue additional recovery abort operations if they are unable to determine in a simple manner whether the state of an FCP I/O operation is ambiguous.

For a target or initiator FCP_Port, an exchange is in an ambiguous state if the FCP_Port has sequence initiative and there exists an unacknowledged frame for the sequence or if the FCP_Port has transferred sequence initiative but the transfer of the initiative has not been confirmed. For a target FCP_Port, an exchange is also in an ambiguous state if the exchange exists between the target FCP_Port and an initiator other than the initiator FCP_Port that performed the TARGET RESET.

2.5 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.

Done.

As part of the FCP control mode page definition, the ECRN bit is defined. See 2.3.

The following significantly modified text is installed in clause 7.1.2 as part of the description of the FCP_CNTL field. This actually changes the numbering, since CRN is the first sub-clause.

7.1.2.1 Command Reference Number (CRN)

The Command Reference Number is provided by the initiator to assist in performing the precise ordering function for FCP commands. If precise ordering is enabled, a non-zero value of CRN shall be treated as a command reference number in determining the receipt and ordering of commands from a particular initiator to the particular logical unit as described in XXX. If precise ordering is enabled, a zero value of CRN shall be ignored and that command shall not be verified for precise ordering. If precise command ordering is not enabled, the value of CRN shall be ignored by the device server. If the FCP_CMND IU specifies a task management function, the CRN shall be ignored and the FCP_CMND IU shall not be verified for precise ordering.

The following significantly modified text is installed after clause 4.2 in a section that becomes the new clause 4.3. This provides the FCP functional model for precise ordering.

4.3 Precise ordering of SCSI commands

In many cases, SCSI communications between an application client and a device server are stateless. In such applications, the precise order of receipt and execution of SCSI commands is often not critical. Any changes in execution order caused by link failures or switch latencies are not important and the recovery and retry mechanisms can be executed while other activities are continued by the application client and the device server.

For those special cases where precise ordering of the receipt of SCSI commands is necessary for the proper operation and error recovery of a device server, FCP-2 defines an additional optional function called precise ordering. An application client can determine if a device server supports the precise ordering function by using the MODE SENSE and MODE SELECT commands to examine and set the enable command reference number (ECRN) bit in the Fibre Channel Control page. See “11.1.1” on page 38.

Precise ordering of SCSI commands uses the COMMAND REFERENCE NUMBER (CRN) in the FCP_CNTL field of the FCP_CMND IU. For each device server having the ECRN bit set to one, the application client places a monotonically increasing one byte integer in the CRN field for each command that is transmitted that also requires precise ordering. The integer begins with a value of one after any Target Reset, LUN Reset, or Fibre Channel Login or Process Login occurs. After the number of precisely ordered commands causes the integer to increment to 255, the integer will wrap back to a value of one. The value of zero is reserved and shall be used for those commands that do not require precise ordering and for task management functions.

The following rules specify how the application client and device server use the CRN to determine that each precisely ordered command has been properly received and executed.

1) A PRLI, Target Reset task management function, and LUN Reset task management function shall reset the CRN to be transmitted by the initiator to 1 and the CRN expected by the device server to 1.
2) The CRN shall be equal to 1 for the first precisely ordered FCP_CMND between the application client and device server and shall be continuously increasing for each subsequent precisely ordered command.
3) 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).
4) The initiator shall not transmit the same CRN again until delivery of the first FCP_CMND transmitted with that CRN has been confirmed by receipt of a FCP_XFER_RDY IU, a FCP_DATA IU, a FCP_RSP IU, an ACK, or an REC response.
5) The device server shall not accept a command with a non-zero CRN into its execution queue until all commands with a previous CRN have been received by the device server.
6) The device server shall accept any command with a CRN having a value of zero into its execution queue and execute the command in accordance with the rules of command queueing. See SAM-2.
7) Task management functions shall have the CRN set to the reserved value of zero and shall not be tested for precise ordering by the device server.

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 precise ordering is not required for that command.

2.6 Confirmed completion of FCP-2 SCSI commands

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 annex C has provided temporary preliminary documentation of the function. That is split out and installed in FCP-2 in the following manner.

The descriptive model is provided as a sub-clause of clause 4.

4.4 Confirmed completion of FCP-2 SCSI commands

Some devices require an acknowledgment of successful delivery of FCP_RSP information. Such an acknowledgment is provided by the confirmed completion function, optionally implemented by FCP-2 devices. The confirmed completion function allows the retry of unsuccessful notifications of errors and confirms that the initiator and the target both agree upon the state of a state dependent device. Retry mechanisms are defined in FC-PH and in profile documents, including FC-PLDA, FC-FLA, and FC-TAPE.

The confirmed completion function may be used to confirm that a SCSI initiator has received an FCP_RSP reporting a SCSI CHECK CONDITION status, together with accompanying sense information. The SCSI target requests in an FCP_RSP IU containing CHECK CONDITION status and sense information that an FCP_CONF be returned by the Initiator. Upon receiving the FCP_CONF, the SCSI target can be assured that the initiator has the information necessary to perform stateful recovery and can then discard its own copy of the information. If the FCP_CONF is not returned, the SCSI target can retransmit the information, assuring eventual receipt of the critical information by the initiator.

The confirmed completion function may be used to confirm that a queued SCSI command has been completed and that the completion information has been successfully transferred to the initiator. The SCSI target requests in an FCP_RSP IU that an FCP_CONF be returned by the initiator. That allows subsequent queued stateful operations to be performed, since the FCP_CONF confirms that the FCP_RSP of GOOD status has been received by the initiator. If the FCP_CONF is not returned, the SCSI target can retransmit the status information, assuring proper synchronization of the state of operations on the initiator and target.

If confirmed completion is enabled, the FCP_CONF may optionally be requested by the FCP_RSP IU for queued commands and for responses containing CHECK CONDITION information. The target may elect not to request FCP_CONF for commands that are stateless, including commands commonly used during initialization (TEST UNIT READY, INQUIRY, and similar commands), even if they are queued or contain CHECK CONDITION information. The target shall not request FCP_CONF for non-queued commands presenting GOOD status and no response information.

If confirmed completion is not enabled, FCP_CONF shall not be requested by the FCP_RSP IU.

Confirmed completion shall not be used by any command linked to another command.

The enabling operation for confirmed completion uses a new bit in PRLI, specified in table in 6.2.6. The bit description is inserted after section 6.2.6.6.

6.2.6.7 Word 3, Bit 7: CONFIRMED COMPLETION ALLOWED

When this bit is set to 1, the process defined by the page is indicating that it supports the confirmed completion protocol for both its target function and its initiator function. If the responder does not support confirmed completion, it shall set this bit to 0 in the PRLI accept page and confirmed completion shall not be used between the logged in processes. When this bit is set to 0, the process defined by the page is indicating that it does not support the confirmed completion protocol and that confirmed completion shall not be used between the logged in processes.

The IU for confirmed completion is specified in the table in section 5.3. as T12. The I5 IU is used to request the confirmation. The same IU is also used to transfer status for intermediate linked commands, but different bits are set in the IU. Appropriate notes are added to the tables.

The IU is defined in a new section installed as section 7.5

7.5 FCP_CONF IU

The FCP_CONF IU has no payload. It is used as described in 4.4 for an initiator to confirm the receipt of the FCP_RSP IU from a target. The frame shall be transmitted by an initiator when the confirmed completion protocol is supported by both the target and the initiator and when the confirmation has been requested by the FCP_CONF_REQ bit in the FCP_RSP.

The FCP_CONF_REQ bit is defined in section 7.4.1 describing the FCP_STATUS byte.

FCP_CONF_REQ, when 1, indicates that the initiator shall transmit an FCP_CONF IU to confirm receipt of the FCP_RSP sequence.

2.7 TERMINATE TASK removal

At present, the TERMINATE TASK function is still included in FCP. With its promotion to obsolete in SPI-3, it will be made obsolete in FCP-2.

2.8 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. Since this proposal has been accepted for persistent reservations in SPC-2 as defined in T10/98-203r7 and above, it is now included in FCP-2.

The wording in SPC-2 is:

For those protocols for which the initiator port’s world wide identification is available to the device server the initiator port’s world wide identification shall be used to determine if the initiator identifier has changed. This determination shall be made at any time the target detects that the configuration of the system may have changed. If the initiator identifier changed, the device server shall assign the new initiator identifier to the existing registration and reservation to the initiator port having the same world wide identification.

The wording in FCP-2 is placed in a new section 5.3 as follows:

5.3 Use of World Wide Names

According to FC-PH, each Fibre Channel node and each Fibre Channel port shall have a world-wide unique name using one of the formats defined by FC-PH and its extensions. Each target and its associated logical units has knowledge of the world-wide unique name of the initiator through the Fibre Channel login process. As a result, the relationship between initiator S_ID and a persistent reservation for a logical unit can be adjusted as defined in SPC-2 during those reconfiguration events that may change the S_ID of the initiator. Any time the target detects that the configuration of the system may have changed and that the initiator identifier changed, the device server shall assign the new initiator identifier to the existing registration and reservation to the initiator port having the same world wide identification.

Each logical unit shall be able to present a world-wide unique name through the INQUIRY command vital product data device identification page (page code 83h) as defined by SPC-2. For FCP devices with a single LUN, the world-wide unique name of the LUN may be the same as the world-wide unique name for the Fibre Channel node.

2.9 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.

(Not yet completed)

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 the flag bit out of SCSI, starting with FCP-2 and SAM-2.

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. If this is approved, I would put it in bytes 6 and 7 of page 19h in section 8.1.1 of the FCP-2 document.

3.3 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.

Should this be included in FCP-2?

4.0 Items not planned for inclusion

4.1 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.

4.2 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. This text is not changed in FCP-2.

4.3 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. This text is not changed in FCP-2.

5.0 Items requiring further discussion

5.1 ECRN bit to determine support of precise ordering

At present, I have treated the ECRN bit as one of those parameters that is always accepted, but ignored if not implemented. That means that the protocol for determining if ECRN is supported is to examine the changeable state of ECRN with MODE SENSE. If it is non-zero, ECRN can be changed. If it the changeable state of ECRN is zero, any attempt to change ECRN with MODE SELECT will be ignored, but no error will be generated.

The CRNS bit proposed for PRLI is not placed in FCP-2, since it is both unnecessary and because the ECRN bit provides all the required architectural functionality in a better place.