



BROCADE

August 31, 2000

T11/00-284 version 5
T10/00-230 revision 5

To: Kumar Malavalli, chairperson, T11, John Lohmeyer, chairperson, T10
 From: Bob Snively
 Date: August 31, 2000
 Subject: FCP-2 items that need to be included in FC-FS

During the creation of FCP-2, a number of items were identified that must be corrected or added to FC-FS. As comment resolution continues, there may be more such requirements and the present requirements may change slightly. This document is intended to track those requirements that have been proposed and those that have been agreed to for FCP-2 and FC-FS.

Changes to FC-FS identified during FCP-2 development

Ref	FCP-2 clause	FCP-2 comment	FC-FS clauses	Proposed	Revision	Approval for FC-FS
1	B.2.1	00-150r3, 1.49	12.6, 12.12, 15.2.2,	05/04/00		June, 2000
2	B.3.1	00-150r3, 1.49	15.4	05/04/00		June, 2000
3	All	00-150r3, 4.18	All	05/10/00		Rejected
4	All	00-150r3, 4.36, 4.96, 4.78, 4.157, 4.175, 4.261	All	05/22/00		Rejected
5	12	00-150r6, 11.2	12.1.1	6/30/00		Rejected
6	8.1	N.A.	15.5.7	7/11/00		7/11/00
7	N.A.	N.A.	15.7, 32.5	7/16/00		Withdrawn
8	8.x	N.A.	15.5.18	7/11/00		? 7/11/00
9		11.21	3.1.45, 12.2, 12.12, 21.6	7/16/00		
10	8.1	11.2, 11.21	15.7	7/16/00		8/xx/00

1 Change requested to ABTS in Annex B.2.1 of FCP-2

The text of B.2.1 is intended to improve the capabilities of the ABTS basic link service. FCP-2 explains the required change with the following text, slightly modified from revision 4:

Brocade Communications Systems, Inc.
 1901 Guadalupe Parkway n San Jose, CA 95131
 T 408.487.8000 n F 408.487.8101
 www.brocade.com

The Abort Sequence Basic Link Service frame shall be used by the Sequence Initiator or Sequence Recipient to request that the ABTS Recipient abort one or more sequences, or abort the entire Exchange, based on the value of bit 0 in the parameter field

ABTS Parameter Field Bit Definitions

Parameter Field	Meaning
Bit 0 = 0	Abort Exchange
Bit 0 = 1	Abort Sequence

To meet the requirements of the FCP-2 standard, the default value of bit 0 in the ABTS request parameter field shall be interpreted as requiring the aborting of the exchange, as described in section 9.1.2.2 on page 41 of FCP-2, revision 4. A value of 1 in bit 0 of the parameter field requires that the sequence be aborted as described in FC-FS and as described in section 12.4 on page 67 of FCP-2 revision 4.

To install this, the following changes must be placed in FC-FS, using T11/99-145v1 as the reference.

In 12.6, table 22, the following modification must be made in the control field definition:

Control Field	Word 2, Bits	Description	Reference
Relative Offset present	3	0 = <u>Parameter field defined for some frames</u> 1 = Parameter Field = Relative Offset	15.2.2

In section 12.6, explanation of word 2 bit 3, the following modification must be made:

Bit 3- Relative Offset present

When bit 3 is set to zero on a Data frame, the Parameter Field is not meaningful. That is, it may be set by the Sequence Initiator, but it shall be ignored by the Sequence Recipient. When bit 3 is set to one in a Data frame, the Parameter Field contains the Relative Offset for the Payload of the frame. Bit 3 is only meaningful on Data frames of a Sequence and shall be ignored on ACK and Link_Response frames. Regardless of the state of bit 3, the parameter field has defined meanings for certain Link Control frames and Basic Link Data frames.

NOTE - When bit 3 is set to 0 on a Data frame, although the Sequence Recipient ignores the value in the Parameter Field, it may pass it to an upper level.

In 12.12, the following changes must be made to reflect the expanded capability demanded of the parameter field.

12.12 Parameter

The Parameter field (Word 5, Bits 31-0) has two meanings based on frame type. For Basic Link Data frames and for Link_Control frames, the Parameter field is used to carry information specific to the individual Link_Control frame. For Data frames, the Parameter field specifies Relative Offset, a four-byte field that contains the relative displacement of the first byte of the

Payload of the frame from the base address as specified by the ULP. Relative Offset is expressed in terms of bytes (see 11.7). The use of the Relative Offset field is optional and is indicated as a Login Service Parameter. The setting of F_CTL bit 3 determines whether the Parameter Field shall be meaningful as a Relative Offset for Data frames.

The offset value shall be relative to an Information Category within a Sequence for an Exchange. If Relative Offset is being used, the number of bytes transmitted in a single Sequence shall not exceed the maximum value of the Relative Offset (Parameter) field.

NOTE - Performance may be improved if data is aligned on natural boundaries.

See clause 21 for a discussion concerning Relative Offset. See clause 14 for a discussion concerning use of the Parameter field in Link_Control frames. [See clause 12 for a discussion concerning use of the Parameter field in Basic Link Data frames.](#)

In 15.2.2, the following change must be made.

15.2.2 Abort Sequence (ABTS)

The Abort Sequence Basic Link Service frame shall be used

- 1) by the Sequence Initiator to request that the Sequence Recipient abort one or more Sequences (see 15.2.2.1 and 23.6.1.1).
- 2) by the Sequence Initiator or Sequence Recipient to request that the ABTS Recipient abort the entire Exchange (see 15.2.2.2).

The decision to attempt to abort one or more Sequences may be determined by the Sequence Initiator (Sequence time-out) or the Sequence Recipient (ACK frame Abort Sequence Condition bits 5-4 or P_RJT frame).

[The Sequence Initiator may optionally require that the Sequence Recipient abort one or more sequences by setting bit 0 of the parameter field to 1. If the parameter field bit 0 is set to 0, the Sequence Recipient may elect to abort one or more sequences or elect to abort the entire Exchange in a protocol specific manner.](#)

The Sequence Recipient may request that one or more Sequences in progress be aborted by setting the Abort Sequence Condition bits to a value of 0 1 on an ACK frame (see 12.6). The ABTS frame may be transmitted without regard to which N_Port holds, or may hold, the Sequence Initiative.

2 Addition of REC to ELS definitions in FC-FS

The following text, presently included as normative clause B.3.1, defines the REC ELS. This is intended to be defined in the standard format as specified in section 15.4 of FC-FS. The text should be placed in 15.4.x, probably (by alphabetical order) just after section 15.4.9. Changes that need to be included are:

On page 169 of FC-FS, table 50 needs an additional line having the following format:

Table 50 - LS_Command codes

Encoded Value (Bits 31-24)	Description	Abbr.
0001 0011	Read Exchange Concise	REC

Just after section 15.4.9, the following description of REC must be installed. This is all new text. The text has been partially updated to reflect the present state of the FCP-2 comment resolution document and the change in context from FCP-2 to FC-FS.

15.4.x Read Exchange Concise (REC)

The REC (Read Exchange Concise) Extended Link Service request Sequence requests an N_Port to return Exchange information for the RX_ID and OX_ID originated by the S_ID specified in the Payload of the request Sequence. The specification of OX_ID and RX_ID shall be provided for the destination N_Port to locate the status information requested. The RX_ID may have the unassigned value of FFFFh if no sequences have been transmitted from the responder or if RX_IDs are not used by the responder. A Responder destination N_Port would use the RX_ID and verify that the OX_ID is consistent, unless the RX_ID was undetermined (i.e., RX_ID = FFFFh). If the RX_ID is specified as undetermined in the request, the Responder must identify the Exchange by means of the S_ID and OX_ID. An Originator N_Port would use the OX_ID and verify that the RX_ID is consistent.

If the destination N_Port of the REC request determines that the Originator S_ID, OX_ID, or RX_ID are inconsistent, then it shall reply with an LS_RJT Sequence with a reason code hex '09' (i.e. Unable to perform command request).

Refer to FC-FS for a description of other applicable LS_RJT reason codes and reason code explanations if the N_Port or SCSI Target is unable to perform the REC request.

The REC shall be sent in a new Exchange.

Protocol:

- Read Exchange Concise request Sequence
- Accept (ACC) reply Sequence

Format:

- FT_1

Addressing:

The S_ID field designates the source N_Port requesting the Exchange information. The D_ID field designates the destination N_Port to which the request is being made.

Payload:

The format of the Payload is shown in the following table.

REC Payload

REC Payload	
Item	Size - Bytes
hex '13000000'	4
Reserved	1
Exchange Originator S_ID	3
OX_ID	2
RX_ID	2

Reply Link Service Sequence:

Service Reject (LS_RJT)

Signifies rejection of the REC request.

ACC

Signifies that the N_Port has transmitted the requested data.

Accept payload:

- The format of the Accept Payload is shown in the table below. The format of the Concise Exchange Status is specified below.

REC Accept Payload

REC Accept Payload	
Item	Size - Bytes
hex '02000000'	4
Concise Exchange Status	
OX_ID	2
RX_ID	2
Originator Address Identifier	4
Responder Address Identifier	4
Data Transfer Count	4
E_STAT	4

The E_STAT is as defined for the Exchange Status Block. The bits specifying whether the Exchange is complete (Bit 29) and whether the responder holds Sequence Initiative (Bit 30) must be valid; the setting of other bits is not required.

The Originator Address Identifier is set to the S_ID with the high-order byte reserved.

The Responder Address Identifier is set to the D_ID with the high-order byte reserved.

The Data Transfer Count is the number of bytes received by the Device Server for a write or the number of bytes transmitted by the Target for a read.

3 Capitalization conventions

Throughout FC-FS, there are grave inconsistencies in capitalization conventions. The following conventions should be used for all entries in FC-FS:

The word “information unit” should always be in lower case. When it is abbreviated, the letters should be all upper case: “IU”. (See T10/00-150r3, comment resolution 4.18.)

This proposal was rejected in the June, 2000 meetings.

4 Capitalization conventions violating modern documentation standards

Throughout FC-FS, many words are capitalized for no particular reason and should be in lower case. The following words should be changed as shown below:

Exchange, Sequence, Originator, and Responder should not be capitalized. (See T10/00-150r3, comment resolution 4.36, 4.96)

Process Login should not be capitalized. (See T10/00-150r3, comment resolution 4.78, 4.157).

Process Associator should not be capitalized. (See T10/00-150r4, comment resolution 4.175).

Link Services, Sequences, Information Category, Unable, Relative Offset, Target, Exchange, Payload, Recipient, Vendor Specific should not be capitalized. (See T10/00-s150r4, comment resolution 4.261)

This proposal was rejected in the June, 2000 meetings.

5 Complete definition of Recovery Qualifier

For historical reasons, the Recovery Qualifier has been defined as:

When one or more Sequences are aborted using the Abort Sequence Protocol (see 23.7.1.1), a Recovery_Qualifier range is identified by the Sequence Recipient which consists of S_ID, D_ID, OX_ID, RX_ID in combination with a range of SEQ_CNT values (low and high). In Class 2 and 3, the Recovery_Qualifier range shall be used by the Sequence Initiator to discard ACK and Link_Response frames and by the Sequence Recipient to discard Data frames.

The work being done in FCP-2 has indicated that recovery capabilities for all protocols would be significantly enhanced if the definition was modified as follows:

When one or more Sequences are aborted using the Abort Sequence Protocol (see 23.7.1.1), a Recovery_Qualifier range is identified by the Sequence Recipient which consists of S_ID, D_ID, OX_ID, RX_ID, and SEQ_ID in combination with a range of SEQ_CNT values (low and high). In Class 2 and 3, the Recovery_Qualifier range shall be used by the Sequence Initiator to discard ACK and Link_Response frames and by the Sequence Recipient to discard Data frames.

This was discovered during the solution of problem 11.2 in document T10/00-150r6.

The proposal was rejected in the July 11, 2000 meeting. There were two major reasons for rejecting this proposal. The first, and most important, was that the change could have unknown and unknowable side effects throughout the document and on present implementations. The second was that there were likely to be multiple SEQ_ID values for which an RRQ had to be established, requiring an even more complex abort procedure.

6 SRR moved from FC-FS to FCP-2.

An earlier correction between revision 3 and revision 4 of FCP-2 had moved SRR from an ELS specified in FC-FS to an FC-4 Link Service specified in FCP-2. As a result, 15.5.7 in FC-FS should be deleted. SRR should also be deleted from table 51 on pdf page 172, from table 187 on pdf page 262, and from table 189 on pdf page 266 (seven occurrences).

This correction was approved in the July 11, 2000 meeting of the FC-FS working group.

7 Correction of references to other FC-4 link services

After review of the text in 15.7 of FC-FS and after inputs from the committee, this proposal is withdrawn. There are presently no FC-4 link services defined in FC-FS.

8 REC moved from FC-FS to FCP-2

This was discussed and approved at the July 11, 2000 meeting of the FC-FS working group, but the approval only involved 6 voting members, so it should be re-approved in the August, 2000 FC-FS working group meeting.

REC should be changed from an ELS to an FC-4 Link Service for two reasons. The first and most important is that the request payload defines an FCP-2 specific use of the concept of relative offset, involving the offset into a multi-sequence string of characters defined in an FCP-2 specific manner. The second is that improvements in the error recovery mechanisms may require FCP-2 specific modifications to the payloads to avoid some ambiguous recovery states. If an REC function is still required in FC-FS, it should be redefined as a separate function with a new name.

The definition of REC should be moved from FC-FS, chapter 15.5.18 and placed in FCP-2 in the proper location in clause 8. Appropriate entries should also be removed from the FC-FS tables referenced in clause 6 of this proposal.

The reference table for FC-4 Link Services proposed in clause 7 of this proposal should also contain REC and the proper reference to FCP-2 provided.

Requires changes if REC is moved.

9 Clarify content of parameter field

When first reviewing the definition of the parameter field and how it could be used as a supplemental FCP-2 command “handle”, the statements:

When bit 3 is set to zero on a Data frame, the Parameter Field is not meaningful. That is, it may be set by the Sequence Initiator, but it shall be ignored by the Sequence Recipient. When bit 3 is set to one in a Data frame, the Parameter Field contains the Relative Offset for the Payload of the frame. Bit 3 is only meaningful on Data frames of a Sequence and shall be ignored on ACK and Link_Response frames. Bit 3 is not meaningful on Basic Link_Data frames.

appear to indicate that the parameter field could be used only as a relative offset and only for Device_Data frames of Information Category Solicited Data. The FCP-2 error recovery requires additional category dependent utilization of the parameter field for other than relative offsets. The following are the minimum changes I expect would be required.

A) In section 3.1.45:

3.1.45 Data frame: ~~A frame containing information meant for FC-4/ULP or the Link application. An FC-4 Device Data frame, an FC-4 Video Data frame, or a Link Data frame. (See clause 14.)~~

B) In section 12.6, where the Relative Offset present bit is discussed:

Bit 3- Relative Offset Present

When bit 3 is set to zero on a Data frame, the ~~value of the~~ Parameter Field ~~is FC-4 protocol specific and may depend on the Information Category. is not meaningful. That is, it may be set by the Sequence Initiator, but it shall be ignored by the Sequence Recipient.~~ When bit 3 is set to one in a Data frame, the Parameter Field contains the Relative Offset for the Payload of the frame ~~as defined by the FC-4 protocol.~~ Bit 3 is only meaningful on Data frames of a Sequence and shall be ignored on ~~ACK and Link_Response on all other~~ frames. Bit 3 is not meaningful on ~~Basic Link_Data Link Control or Basic Link Service Link Data~~ frames.

~~NOTE~~—When bit 3 is set to 0 on a Data frame, ~~although the Sequence Recipient ignores~~ the value in the Parameter Field, ~~it may pass it to an~~ shall be passed to the upper level.

D) In section 12.12, where the Parameter field is discussed:

12.12 Parameter

The Parameter field (Word 5, Bits 31-0) has ~~two~~ meanings based on frame type. For Link_Control frames, the Parameter field is used to carry information specific to the individual Link_Control frame.

For Data frames with the Relative Offset Present bit set to 1, the Parameter field specifies Relative Offset, a four-byte field that contains the relative displacement of the first byte of the Payload of the frame from the base address as specified by the ULP. Relative Offset is expressed in terms of bytes (see 11.7). The use of the Relative Offset field is optional and is indicated as a Login Service Parameter. ~~The setting of F_CTL bit 3 determines whether the Parameter Field shall be meaningful as a Relative Offset for Data frames. The offset value shall be relative to an Information Category within a Sequence for an Exchange.~~ If Relative Offset is being used, the number of bytes transmitted relative to the protocol-specific base address in a single Sequence shall ~~not exceed~~ be less than the maximum value of the Relative Offset (Parameter) field (2^{32}).

For Data frames with the Relative Offset Present bit set to 0, the Parameter field shall be set and interpreted in a protocol specific manner, which may depend on the type of Information Unit carried by the frame.

~~NOTE~~— ~~Performance may be improved if data is aligned on natural boundaries.~~

See clause 21 for a discussion concerning Relative Offset. See clause 14 for a discussion concerning use of the Parameter field in Link_Control frames.

E) In section 21.6, table 207, all entries that indicate “Parameter field not meaningful” should be replaced with entries that indicate “Parameter field meaning is protocol and IU specific”.

10 Allow change of sequence initiative with certain FC-4 link services

The error recovery procedures required in FCP-2 require that sequence initiative be recovered when SRR is executed so that the proper recovery events can be executed. At present, the only function defined by FC-FS that can over-ride sequence initiative is ABTS and RSI. The text of FC-FS section 15.7 should be changed as follows to correct this. This was accepted in the August FC-FS working group by a vote of 6 in favor, 3 opposed, and 15 abstaining.

15.7 FC-4 Link Service

A FC-4 Link Service request solicits a destination Port (F_Port or N_Port) to perform a function or service in order to support an individual FC-4 Device_Data protocol. The Information Category for a request shall be specified as Unsolicited Control. A FC-4 Link Service reply may be transmitted in answer to a FC-4 Link Service request. The Information Category for a reply shall be specified as Solicited Control. Each request or reply shall be composed of a single Sequence. The format of the request or reply shall be specified by the individual FC-4 being supported and is outside the scope of FC-FS. Each Sequence may be composed of one or more frames.

The protocols supported by the FC-4 Link Services shall be performed within a single Exchange, intended exclusively for the purpose. FC-4 Link Service protocols are performed using a two Sequence Exchange. The protocols consist of a request Sequence by the Originator (N_Port), transfer of Sequence Initiative, and a reply Sequence from the Responder (N_Port or F_Port). The execution of a FC-4 Link Service may perform sequence abort functions and modify sequence initiative of other exchanges in a protocol specific manner. The Sequence Initiator and Sequence Recipient shall follow the rules for Sequence management and Recovery_Qualifier reuse as specified in clause 18.