

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 22 March 2007
 Subject: 07-072r1 FCP-4 QUERY TASK task management function

Revision history

Revision 0 (14 February 2007) First revision
 Revision 1 (22 March 2007) Incorporated comments from March 2007 FCP WG.

Related documents

sam4r08 - SCSI Architecture Model - 4 (SAM-4) revision 8
 fcp4r00 - Fibre Channel Protocol - 4 (FCP-4) revision 0
 T11/fc-fs-2-r1.01 - Fibre Channel Framing and Signaling - 2 (FC-FS-2) revision 1.01
 T11/fc-ls-r1.62 - Fibre Channel Link Services (FC-LS) revision 1.62
 T11/00-133v1 BA_ACC vs BA_RJT usage (Carl Zeitler, Compaq Computer Corp)
 07-066 - SAM-4 SAS-2 QUERY TASK SET task management function (Rob Elliott, HP)
 07-067 - SAM-4 SAS-2 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)
 07-143 - FCP-4 QUERY TASK SET task management function (Rob Elliott, HP)
 07-144 - FCP-4 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)

Overview

1. FCP-4 doesn't claim to support the SCSI QUERY TASK task management function, although it really does functionally support it. QUERY TASK translates into the Read Exchange Consise (REC) Extended Link Service (ELS) defined in FC-LS, just like ABORT TASK translates into the Abort Sequence (ABTS) Basic Link Service defined in FC-FS-2. QUERY TASK should be described as supported in FCP-4.

2. David Peterson (Brocade) noted on the T11.2 and T10 reflectors that FC-LS and FCP-4 have conflicting requirements for the Reason Code value when a REC ELS needs to be rejected with a Reason Code Explanation of 17h (i.e., Invalid OX_ID-RX_ID combination):

- a) FC-LS specifies the reason code is 09h (i.e., Unable to perform command request);
- b) FCP-4 specifies that the reason code is 03h (i.e., Logical error).

Feedback indicates that both need to be allowed.

3. Miscellaneous changes to ABTS are proposed to update wording to match FC-FS-2 and FC-LS. For example, ABTS-LS is the acronym for ABTS(Exchange).

4. Complete SCSI Service Response mappings are proposed for all the task management functions. FCP-4 does not currently provide this level of detail, although SCSI transport protocols are supposed to do so.

Suggested changes to FCP-4

3.2 Abbreviations

ABTS Abort Sequence [Basic Link Service](#) (see FC-FS-2)

[ABTS-LS ABTS Basic Link Service with the Parameter field bit 0 set to zero \(i.e., Abort Exchange\) \(see FC-FS-2\)](#)

Editor's Note 1: FC-FS-2 and FC-LS use "ABTS-LS" (i.e., Last Sequence) as the abbreviation for an ABTS that performs the Abort Exchange. FCP often uses "ABTS (Exchange)" to denote that use of ABTS. Changes are proposed to use the ABTS-LS term instead. Not all uses are highlighted for change (e.g., there are some uses in the annex C figures); the editor should globally change ABTS (Exchange) to ABTS-LS.

Editor's Note 2: There is no FC-FS-2 and FC-LS acronym to replace "ABTS (Sequence)". Perhaps

[ABTS-NLS should be added \(Not Last Sequence\) and used like ABTS-LS.](#)

...

[BA_ACC Basic Accept Basic Link Service \(see FC-FS-2\)](#)

[BA_RJT Basic Reject Basic Link Service \(see FC-FS-2\)](#)

[REC Read Exchange Concise ELS \(see FC-LS\)](#)

[RRQ Reinstate Recover Qualifier ELS \(see FC-LS\)](#)

...

4 General

4.6 Retransmission of unsuccessfully transmitted IUs

Error detection and IU retransmission algorithms are defined in clause 12.

The Read Exchange Concise (REC) ELS may be used by the initiator FCP_Port to determine the state of an ongoing Exchange. See 6.5.

Support for the REC ELS by both the initiator FCP_Port and target FCP_Port is indicated by the REC_SUPPORT bit in the PRLI request FCP Service Parameter page ([see 6.3.4](#)) and [the](#) PRLI accept FCP Service Parameter page ([see 6.3.4 and 6.3.5](#)).

If the target FCP_Port responds with the REC_SUPPORT bit set to one and an error is identified by any of the detection mechanisms defined in clause 12, then the initiator FCP_Port may use the REC ELS to determine the nature of the error.

Target FCP_Ports that do not support the REC_SUPPORT bit indicate they do not support REC by [performing](#)[returning](#) a Link Service Reject (LS_RJT) [ELS](#) in response to an REC ELS. See 8.3.

4.9 Task management [functions](#)

[4.9.1 Task management functions overview](#)

An application client requests a task management function ([see SAM-4](#)) to control explicitly the processing of one or more FCP I/O operations ([see 9.2.2.5](#)).

The task management function mappings are specified in table 3.

Table 3 — Task management functions, SAM-3 to FCP-4 [\[moved up a few paragraphs\]](#)

SAM-3 task management function	FCP-4 equivalentimplementation
ABORT TASK	ABTS-LS Basic Link Service (see FC-FS-2) FCP recovery abort ^a
ABORT TASK SET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to 02h (i.e., ABORT TASK SET)
CLEAR TASK SET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to 04h (i.e., CLEAR TASK SET)
CLEAR ACA	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to 40h (i.e., CLEAR ACA)
LOGICAL UNIT RESET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to 10h (i.e., LOGICAL UNIT RESET)
QUERY TASK	REC ELS (see FC-LS)

^a FC-FS-2 basic link services are used to perform the ABORT TASK [task management](#) function.

~~The ABORT TASK task management function is mapped to the FC-FS-2 ABTS basic link service while the other task management functions are mapped into control bits (see table 20) in the FCP_CMND IU. Task management functions that use the FCP_CMND IU are transmitted as the first IU in a new Exchange. A task management function that uses the FCP_CMND IU ends with an FCP_RSP IU that indicates the completion status of the function. If the addressed logical unit is not supported or is not available (e.g., not connected or not configured) the FCP_CMND IU target FCP_Port:~~

- a) should end the Exchange with an FCP_RSP IU completion status of 09h (i.e., Task Management function incorrect logical unit number) (see table 24); and
- b) may end the Exchange with an FCP_RSP IU completion status of 00h (i.e., Task Management function complete) (see table 24).

The FCP_CDB field in FCP_CMND IUs that perform task management functions is ignored.

~~The QUERY TASK task management function is not supported.~~

FC-FS-2 basic link services and FC-LS extended link services are used to perform the ABORT TASK task management function, to perform the QUERY TASK task management function, to recover Exchange resources, and to re-establish other initial conditions.

Table 4 defines the SCSI Service Response mappings for an initiator FCP_Port for FCP_CMND delivered task management functions

Table 4 — SCSI Service Response mapping for FCP_CMND IU delivered task management functions

<u>SCSI Service Response</u>	<u>Response to the FCP_CMND IU</u>
<u>FUNCTION COMPLETE</u>	<u>FCP_RSP IU with the RSP_CODE field set to 00h (i.e., Task Management function complete)</u>
<u>FUNCTION REJECTED</u>	<u>FCP_RSP IU with the RSP_CODE field set to 04h (i.e., Task Management function rejected)</u>
<u>INCORRECT LOGICAL UNIT NUMBER</u>	<u>FCP_RSP IU with the RSP_CODE field set to 09h (i.e., Task Management function incorrect logical unit number)</u>
<u>SERVICE DELIVERY OR TARGET FAILURE</u>	<u>All other responses, including:</u> a) <u>timeout; and</u> b) <u>FCP_RSP IU with the RSP_CODE field set to 05h (i.e., Task Management function failed)</u>

4.9.2 ABORT TASK task management function

The ABORT TASK task management function causes the device server to abort the specified task, if it exists, using the ABTS-LS Basic Link Service (see FC-FS-2) recovery abort protocol, if the task exists. ~~The action is defined in SAM-3. The ABORT TASK task management function is performed by the initiator FCP_Port (i.e., Exchange Originator) using the recovery abort (see 12.3).~~

~~The specified Exchange shall be terminated by the initiator FCP_Port using the recovery abort.~~ To be compliant with FC-FS-2, the ABORT TASK task management function may not immediately release all Exchange resources, since a Recovery_Qualifier may be established to allow for the management of information that may already have been delivered to the fabric.

In addition to recovering Exchange resources that may have been left unavailable while processing task management functions, recovery abort (see 12.3) uses the ABTS-LS Basic Link Service may be used to recover Exchange resources left in an undefined state by any of the task abort events defined in SAM-3 or by any similar events.

Table 5 defines the SCSI Service Response mappings for an initiator FCP Port for the ABORT TASK task management function.

Table 5 — SCSI Service Response mapping for ABORT TASK

<u>SCSI Service Response</u>	<u>Response to the ABTS-LS</u>
<u>FUNCTION COMPLETE</u>	Either: a) <u>BA ACC</u> ; or b) <u>BA RJT</u> with Reason Code set to 03h (i.e., Logical error) and Reason Code Explanation set to 03h (i.e., Invalid OX ID-RX ID combination)
<u>FUNCTION REJECTED</u>	<u>BA RJT</u> with Reason Code set to 01h (i.e., Invalid command code) or 09h (i.e., Unable to perform command request)
<u>INCORRECT LOGICAL UNIT NUMBER</u>	<u>No responses map to this SCSI Service Response</u>
<u>SERVICE DELIVERY OR TARGET FAILURE</u>	<u>All other responses, including timeout</u>

4.9.3 QUERY TASK task management function

The QUERY TASK task management function maps to the REC ELS. The REC ELS shall be transmitted in a new Exchange.

Table 6 defines the SCSI Service Response mappings for an initiator FCP Port for the QUERY TASK task management function.

Table 6 — SCSI Service Response mapping for QUERY TASK

<u>SCSI Service Response</u>	<u>Response to the REC ELS</u>
<u>FUNCTION COMPLETE</u>	Either: a) <u>LS ACC</u> with E_STAT (i.e., word 5) bit 29 (i.e., Completion) set to one (i.e., complete); or b) <u>LS RJT</u> with Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to perform command request) and Reason Code Explanation set to 17h (i.e., Invalid OX ID-RX ID combination)
<u>FUNCTION SUCCEEDED</u>	<u>LS ACC</u> with E_STAT (i.e., word 5) bit 29 (i.e., Completion) set to zero (i.e., open)
<u>FUNCTION REJECTED</u>	<u>LS RJT</u> with Reason Code set to: a) <u>01h</u> (i.e., Invalid ELS Command code); or b) <u>0Bh</u> (i.e., Command not supported)
<u>INCORRECT LOGICAL UNIT NUMBER</u>	<u>No responses map to this SCSI Service Response</u>
<u>SERVICE DELIVERY OR TARGET FAILURE</u>	<u>All other responses, including:</u> a) <u>timeout</u> ; and b) <u>LS RJT</u> with Reason Code set to 09h (i.e., Unable to perform command request) and Reason Code Explanation set to 15h (i.e., Invalid Originator S_ID)

6.3.4 PRLI request FCP Service Parameter page format

...

Word 3, Bit 10: REC_SUPPORT: When the REC ELS supported (REC_SUPPORT) bit is set to one, the Originator is indicating that it supports, as an initiator FCP_Port, the transmission of the REC ELS. The capability of the initiator FCP_Port to retransmit unsuccessfully transmitted data is determined by the RETRY bit (i.e., a

REC_SUPPORT bit set to one does not indicate the initiator FCP_Port supports retransmission of data). When the REC_SUPPORT bit is set to zero, the Originator is providing no information about whether it supports transmission of the REC ELS.

...

6.3.5 PRLI accept FCP Service Parameter page format

...

Word 3, Bit 10: REC_SUPPORT: When the REC ELS supported (REC_SUPPORT) bit is set to one, the Responder is indicating that it supports, as a target FCP_Port, the receipt of the REC ELS. The capability of the target FCP_Port to retransmit unsuccessfully transmitted data is determined by the RETRY bit (i.e., a REC_SUPPORT bit set to one does not indicate the target FCP_Port supports retransmission of data). When the REC_SUPPORT bit is set to zero, the Responder is indicating that it may not support receipt of the REC ELS.

6.5 Read Exchange Concise (REC)

See FC-LS for a description of the REC ELS. FCP-4 specific usage of REC is as follows:

- a) if task retry identification is active for the Originator and the Responder, the PARAMETER field of the request Sequence shall contain the task retry identifier for the task specified by the OX_ID field value and RX_ID field value;
- b) if the destination FCP_Port of the REC request determines that the ORIGINATOR S_ID, OX_ID, or RX_ID fields, or task retry identifier are inconsistent, then it shall respond with an LS_RJT Sequence with ~~a reason code of "Logical error" and a reason code explanation of "Invalid OX_ID-RX_ID combination"~~the Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to process command request); and the Reason Code Explanation set to 17h (i.e., Invalid OX_ID-RX_ID combination);
- c) the REC shall be sent in a new Exchange. The Exchange shall be ended by the response to the REC;
- d) if the RX_ID field value in the REC request payload was FFFFh, the RX_ID field value in the REC accept payload may be set to the value selected by the Responder when the first frame of the Exchange was received;

and
- e) the FC4VALUE field is the number of bytes successfully received by the Device Server for a write or the number of bytes transmitted by the target FCP_Port for a read.

9.2 FCP_CMND IU

9.2.1 Overview and format of FCP_CMND IU

9.2.2 FCP_CMND IU field descriptions

9.2.2.5 TASK MANAGEMENT FLAGS field

The TASK MANAGEMENT FLAGS field specifies the task management function to be performed, if any.

Task management functions shall be requested by the initiator FCP_Port (Exchange Originator) using a new Exchange. If the TASK MANAGEMENT FLAGS field is set to a nonzero value, the FCP_CDB field, the FCP_DL field, the TASK ATTRIBUTE field, the RDDATA bit, and the WRDATA bit shall be ignored and the FCP_BIDIRECTIONAL_READ_DL field shall not be included in the FCP_CMND IU payload. If the TASK MANAGEMENT FLAGS field is set to a reserved value, the target FCP_Port shall return an FCP_RSP IU containing the RSP_CODE field set to 02h (i.e., FCP_CMND fields invalid).

The clearing actions performed by task management functions are specified in table 5.

The format of the TASK MANAGEMENT FLAGS field is specified in table 20.

Table 7 — TASK MANAGEMENT FLAGS field

Code	Task management function ^a
00h	None
01h	QUERY TASK SET per 07-143
02h	ABORT TASK SET
04h	CLEAR TASK SET
08h	QUERY UNIT ATTENTION per 07-144
10h	LOGICAL UNIT RESET
20h	Obsolete
40h	CLEAR ACA
80h	Obsolete
All others	Reserved
^a The ABORT TASK and QUERY TASK task management functions are specified in 4.9.	

The **CLEAR ACA** task management function causes the ACA condition to be cleared. When the task manager clears the ACA condition, any task within that task set may be completed subject to the rules for task management specified by SAM-3. If there is no ACA condition present, the CLEAR ACA task management function shall be accepted and the FCP_RSP IU shall contain a RSP_CODE field set to 00h (i.e., Task Management function complete).

The use of the ACA bit in the CDB control field and the implementation of ACA is described in SAM-3.

Depending on the mode page parameters that have been established (see SPC-3), additional FCP I/O operations may have to be aborted by the recovery abort as part of the process of clearing the automatic contingent allegiance.

The CLEAR ACA is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange.

Support of the CLEAR ACA task management function is mandatory in the Fibre Channel Protocol if the FCP device sets the NORMACA bit to one in the INQUIRY data. It shall not be sent to a logical unit with a NORMACA bit equal to zero in the INQUIRY data.

The **LOGICAL UNIT RESET** task management function performs a LOGICAL UNIT RESET task management function as defined in SAM-3. LOGICAL UNIT RESET aborts all tasks in the task set for the logical unit and performs a LOGICAL UNIT RESET for all dependent logical units. Support of the LOGICAL UNIT RESET task management function is mandatory for the Fibre Channel Protocol.

The LOGICAL UNIT RESET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. LOGICAL UNIT RESET resets the internal states of the target FCP_Port and logical unit as shown in 4.10. Exchange resources to be cleared may be cleared by the following mechanisms:

- a) a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the LOGICAL UNIT RESET for each task in the logical unit known to that initiator FCP_Port;
- b) a task, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET is ended in the logical unit. The initiator FCP_Port for that task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or

- c) a task for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-3.

NOTE 3 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a LOGICAL UNIT RESET task management function if the Control mode page indicates that the TASK ABORTED status is supported.

The **CLEAR TASK SET** task management function causes all tasks from all initiator FCP_Ports in the specified task set to be aborted as defined in SAM-3. Support of the CLEAR TASK SET task management function is mandatory for the Fibre Channel Protocol.

The CLEAR TASK SET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. CLEAR TASK SET resets internal states of the target FCP_Port as shown in 4.10. Exchange resources to be cleared may be cleared by one or more of the following mechanisms:

- a) a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the CLEAR TASK SET for each task known to that initiator FCP_Port;
- b) a task, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET is ended in the logical unit. The initiator FCP_Port for that task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or
- c) a task for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-3.

NOTE 4 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a CLEAR TASK SET task management function if the Control mode page indicates that the TASK ABORTED status is supported.

The **ABORT TASK SET** task management function requests the ABORT TASK SET task management function to be performed as defined in SAM-3. Support of the ABORT TASK SET task management function is mandatory in the Fibre Channel Protocol.

The ABORT TASK SET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. ABORT TASK SET resets internal states of the target FCP_Port as shown in 4.10. Exchange resources may be cleared by a recovery abort sequence (see 12.3) generated by the initiator FCP_Port that sent the ABORT TASK SET for each task known to the initiator FCP_Port.

9.2.2.8 FCP_CDB field

The FCP_CDB field contains the CDB to be sent to the addressed logical unit. The maximum CDB length is 16 bytes unless the ADDITIONAL_FCP_CDB_LENGTH field has specified that there is an ADDITIONAL_FCP_CDB field. The FCP_CDB field shall be ignored if the task management flags field is set to a nonzero value.

The CDB format is defined by SAM-3 and the contents of the CDB are defined in the SCSI command standards.

Bytes between the end of a CDB and the end of the FCP_CDB field or, if applicable, the ADDITIONAL_FCP_CDB field shall be reserved.

9.2.2.9 ADDITIONAL_FCP_CDB field

The ADDITIONAL_FCP_CDB field contains any CDB bytes beyond those contained within the 16 byte FCP_CDB field.

The ADDITIONAL_FCP_CDB field shall not be present if the task management flags field is set to a nonzero value. The contents of the field shall be those bytes of an extended CDB beyond the first 16 bytes of the CDB as defined in the SCSI command standards.

9.5 FCP_RSP IU

9.5.16 FCP_RSP_INFO field

The FCP_RSP_INFO field contains information describing only the protocol failures detected during the processing of an FCP I/O operation. If none of the specified protocol failures have occurred, the FCP_RSP_INFO field shall not be included in the FCP_RSP IU and the FCP_RSP_LEN_VALID bit shall be zero. The FCP_RSP_INFO does not contain link error information, since FC-FS-2 provides the mechanisms for presenting such errors. The FCP_RSP_INFO field does not contain SCSI logical unit error information, since that is contained in the FCP_SNS_INFO field as described in 9.5.17. The FCP_RSP_INFO field shall contain valid information if the target FCP_Port detects any of the conditions indicated by an FCP FCP_RSP_CODE. The format of the FCP_RSP_INFO field is specified in table 23.

Table 23 — FCP_RSP_INFO field format

Byte/Bit	7	6	5	4	3	2	1	0
0	Reserved							
2	Reserved							
3	RSP_CODE							
4	Reserved (if any)							
7	Reserved (if any)							

The valid RSP_CODE values are specified in table 24.

Table 24 — RSP_CODE field

Code	Description
00h	Task Management function complete
01h	FCP_DATA length different than FCP_BURST_LEN
02h	FCP_CMND fields invalid
03h	FCP_DATA parameter mismatch with FCP_DATA_RO
04h ^a	Task Management function rejected
05h	Task Management function failed
08h ^a	Task Management function succeeded
09h ^a	Task Management function incorrect logical unit number
06h - 07h 0Ah - FFh	Reserved
^a Only valid when responding to task management functions	

The completion status of the task management function is indicated by the RSP_CODE field. If the Exchange is aborted before the FCP_RSP IU is returned, the completion status is unknown. If the RSP_CODE field is set to 05h (i.e., Task Management function failed), the state of the logical unit is unknown.

Activities started by a task management function may continue after the FCP_RSP IU for the task management has been delivered.

9.5.17 FCP_SNS_INFO field

The FCP_SNS_INFO field contains the autosense data specified by SPC-3. The proper FCP_SNS_INFO shall be presented when the SCSI status byte of CHECK CONDITION is presented as specified by SAM-3. If no conditions requiring the presentation of SCSI sense data have occurred, the FCP_SNS_INFO field shall not be

included in the FCP_RSP IU and the FCP_SNS_LEN_VALID bit shall be zero. FCP devices shall perform autosense.

11.5 Read Exchange Concise Timeout Value (REC_TOV)

REC_TOV is used by the initiator FCP_Port to provide a minimum polling interval for REC and by the target FCP_Port for FCP_CONF IU error detection. The REC_TOV timer shall be implemented such that at least one REC_TOV period passes between transmission of a command and the first polling for Exchange status with the REC ELS. Table 31 describes REC_TOV usage pertaining to the initiator FCP_Port.

...

12.3.2 Initiator FCP_Port invocation of recovery abort

The initiator FCP_Port terminating the Exchange sends an ABTS-LS ~~sequence~~ to the D_ID of the target FCP_Port of the Exchange being terminated. ~~The ABTS sequence shall have the PARAMETER field set to ABORT EXCHANGE.~~ The ABTS-LS ~~sequence~~ shall be generated using the OX_ID and RX_ID field values of the Exchange to be aborted. FC-FS-2 allows ABTS-LS to be generated by an FCP_Port regardless of whether or not it has Sequence Initiative. Following the transmission of ABTS-LS, any Device_Data Frames received for this Exchange shall be discarded until the BA_ACC with ~~“Last Sequence of Exchange” bit set to one~~ the F_CTL field Last Sequence bit set to one (i.e., last Sequence of the Exchange) is received from the target FCP_Port.

Recovery abort may not take effect immediately (e.g., if ABTS-LS is sent following transmission of a READ command, the initiator FCP_Port may receive some or all of the requested read data before receiving the BA_ACC ~~for~~ the ABTS). The initiator FCP_Port shall be capable of receiving this data and providing BB_Credit in order for the target FCP_Port to send the BA_ACC.

After the processing of any task management function that clears tasks (e.g., ABORT TASK SET, CLEAR TASK SET, and LOGICAL UNIT RESET), recovery abort shall be invoked for all Exchanges not successfully terminated with an FCP_RSP IU status set to COMMAND CLEARED. (See 9.2.2.5).

Following receipt of the BA_ACC in response to an ABTS-LS, and after R_A_TOV_SEQ_QUAL has elapsed, the initiator FCP_Port shall transmit an RRQ ELS (see FC-LS).

If a BA_ACC, BA_RJT, LOGO, or PRLO is not received from the target FCP_Port within 2 times R_A_TOV_ELS, second level error recovery as described in 12.5 shall be performed.

12.3.3 Target FCP_Port response to recovery abort

When an ~~ABTS (Abort Exchange)~~ ABTS-LS is received at the target FCP_Port, it shall abort the designated Exchange and return one of the following responses:

- a) the target FCP_Port shall discard the ABTS-LS and return LOGO if the N_Port or NL_Port issuing the ABTS-LS is not currently logged in (i.e., no PLOGI);
- b) the target FCP_Port shall return BA_RJT with ~~Last Sequence of Exchange bit set to one~~ the F_CTL field Last Sequence bit set to one (i.e., last Sequence of the Exchange) if the received ABTS-LS contains an assigned RX_ID field value and a FQXID that is unknown to the target FCP_Port; or
- c) the target FCP_Port shall return BA_ACC with ~~Last Sequence of Exchange bit set to one~~ the F_CTL field Last Sequence bit set to one (i.e., last Sequence of the Exchange).

Upon transmission of any of the above responses, the target FCP_Port may reclaim any resources associated with the designated Exchange after R_A_TOV_SEQ_QUAL has elapsed or a Reinstate Recovery Qualifier (RRQ) ELS request has been received.

If the RX_ID field is FFFFh, target FCP_Ports shall qualify the FQXID of the ABTS-LS based only upon the combined values of the D_ID field, S_ID field, and the OX_ID field, not the RX_ID field.

If the Exchange resources were not reclaimed upon responding to the ABTS-LS, they shall be reclaimed at the time the response to the RRQ is sent.

When an RRQ is received at the target FCP_Port, it shall return one of the following responses:

- a) the target FCP_Port shall discard the RRQ and return LOGO if the N_Port or NL_Port issuing the RRQ is not currently logged in (i.e., no PLOGI);

- b) ~~if the received RRQ contains an RX_ID field value, other than FFFFh, that is unknown to the target FCP_Port, then~~ the target FCP_Port shall return LS_RJT with ~~Last Sequence of Exchange bit set to one if the received RRQ contains an RX_ID field value, other than FFFFh, that is unknown to the target FCP_Port. The reason code shall be "Logical error" with a reason code explanation set to "Invalid OX_ID-RX_ID combination"~~
- A) the F_CTL field Last Sequence bit set to one (i.e., last Sequence of the Exchange);
B) the Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to process command request); and
C) the Reason Code Explanation set to 17h (i.e., Invalid OX_ID-RX_ID combination);
- ; or
- c) the target FCP_Port shall return ACC with Last Sequence of Exchange bit set to one.

12.3.4 Additional error recovery by initiator FCP_Port

This procedure may be used whether or not the FCP devices have agreed to Sequence level recovery.

If ULP_TOV times out and the Exchange is not complete, the application client shall clear the Exchange resources using the ABORT TASK task management request or the initiator FCP_Port shall clear the Exchange resources using the recovery abort protocol (see 4.9).

12.3.5 Additional error recovery by target FCP_Port

This procedure may be used whether or not the FCP devices have agreed to Sequence level recovery.

If a target FCP_Port detects a Sequence error, it shall discard the Sequence(s) based on the Exchange error policy specified by the F_CTL field Abort Sequence Condition bits in the first frame of the Exchange (see FC-FS-2).

For acknowledged classes of service, if a target FCP_Port detects a Sequence error, it may abort the sequence by sending an ABTS ~~with the PARAMETER field to ABORT SEQUENCE~~ with the Parameter field bit 0 set to one (i.e., Abort Sequence) (see FC-FS-2). If a Recovery Qualifier range is returned in the BA_ACC for the ABTS the target FCP_Port shall send a RRQ ELS after R_A_TOV_SEQ_QUAL times out after receipt of the BA_ACC.

For unacknowledged classes of service, the target FCP_Port shall not attempt recovery for Sequence errors. The target FCP_Port shall depend on initiator FCP_Port timeouts for recovery.

Target FCP_Ports shall implement RR_TOV as described in 11.4 to facilitate recovery of resources allocated to an initiator FCP_Port that is no longer responding. The target FCP_Port may send a LOGO to the initiator FCP_Port and terminate all open Exchanges for that initiator FCP_Port upon detection of the following:

- a) the initiator FCP_Port has failed to perform target FCP_Port Exchange authentication within RR_TOV_AUTH (see FC-DA); or
- b) RR_TOV_SEQ_INIT times out without the initiator FCP_Port transmitting any expected Sequence for any open Exchange at this target FCP_Port (e.g., FCP write Data-In response to an FCP_XFER_RDY IU).

12.4.2 Additional error recovery requirements

12.4.2.1 Error indicated in ACK

If an ACK is received with the F_CTL field Abort Sequence Condition bits set to Abort Sequence, Perform ABTS, the Sequence Initiator shall send an ABTS for the Sequence. After R_A_TOV times out, an RRQ shall be sent by the Sequence Initiator.

12.4.2.2 Missing ACK

FC-FS-2 requires that an ABTS(Sequence) be transmitted by a Sequence Initiator detecting a missing ACK. If no ACK has been received within E_D_TOV, the target FCP_Port shall abort the sequence by sending an ABTS request with the PARAMETER field set to ABORT SEQUENCE. If a Recovery Qualifier range is returned in the BA_ACC for the ABTS the target FCP_Port shall send an RRQ at least R_A_TOV_SEQ_QUAL after receipt of the BA_ACC. Adjustment of subsequent sequence counts may be required as specified by FC-FS-2.

12.5.2 REC error recovery

If a response to an REC is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port shall:

- 1) send an ~~ABTS(Exchange)~~ABTS-LS for the REC followed by an RRQ if a BA_ACC is received for the ABTS-LS; and
- 2) send another REC in a new Exchange.

If the response to the second REC is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port should send an ~~ABTS(Exchange)~~ABTS-LS for the REC followed by an RRQ if a BA_ACC is received for the ABTS-LS;

Other retry mechanisms after the second REC fails are optional and, if implemented, shall comply with FC-FS-2.

~~ABTS(Exchange)~~ABTS-LS may be required to clear resources associated with the original failing Exchange if the retry mechanisms are not successful.

See figure C.26 through figure C.29.

12.5.3 SRR error recovery

If a response to an SRR is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port shall send an ~~ABTS(Exchange)~~ABTS-LS for the SRR followed by an RRQ if a BA_ACC is received for the ABTS-LS. The initiator FCP_Port shall then perform an ~~ABTS(Exchange)~~ABTS-LS for the original Exchange.

See figure C.30 through figure C.33.