

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 17 November 2007
 Subject: 07-519r0 FCP-4 Set REC FC4VALUE to zero for bidirectional commands

Revision history

Revision 0 (17 November 2007) First revision

Related documents

fcp4r00 - Fibre Channel Protocol - 4 (FCP-4) revision 0
 T11/fc-ls-1.62 - Fibre Channel Link Services (FC-LS) revision 1.62

Overview

The FCP-4 description of the REC LS_ACC payload FC4VALUE field specifies that it be set to “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.”

1. This doesn't specify how to handle bidirectional commands, which transfer both read and write data. Since the FC4VALUE field is only four bytes long and there are only two Reserved bytes remaining in the payload, there is not room for two full four-byte values. So, a value of 00000000h is proposed.
2. FCP-4 should mention that non-data commands (e.g., TEST UNIT READY) report 0000000h.
3. Since the REC ELS can be used to query about any Exchange, FCP-4 should also specify that 00000000h be reported for a task management function (since they're defined in FCP-4), an Exchange for an SRR ELS (since it is defined in FCP-4), and an Exchange for another REC ELS (not defined by FCP-4, but used by it...note that FCP-4 doesn't ever mention sending a REC ELS to query about another one).
4. If a read or write command performs retransmission (due to sequence-level recovery) or overlay (per the EMDP bit in the Disconnect-Reconnect mode page and process login), it is unclear how to report a single value. Wording from 9.4.2 is copied that says such data should be “counted only once.”

Suggested changes to FCP-4

6.5 Read Exchange Concise (REC)

See FC-LS for a description of the REC ELS. FCP-4 specific usage of REC [ELS](#) 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”;
- 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~~LS_ACC 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~~in the REC LS_ACC payload shall be set to:
 - A) for a write command, the number of bytes successfully received by the device server. Data that has been retransmitted or overlaid shall be counted only once;
 - B) for a read command, the number of bytes transmitted by the target FCP_Port. Data that has been retransmitted or overlaid shall be counted only once;
 - C) for a non-data command, 00000000h; and
 - D) for a bidirectional SCSI command, 00000000h; and
 - E) for any other Exchange defined by this standard (e.g. a task management function, an SRR ELS, or another REC ELS), 00000000h.

9.4.2 FCP_DATA IUs for read and write operations

...

If the amount of data requested or transferred does not match the number of bytes of data calculated from the value of the FCP_DL field and the value of the FCP_RESID field (see 9.5.12), the error detection and recovery procedure described in clause 12 may be invoked or the FCP I/O operation may be terminated with a recovery abort (see 12.3) or other failure indication. The mechanism an initiator FCP_Port uses to determine that the correct amount of data has been returned is vendor specific. Data that has been retransmitted and overlaid shall be counted only once for the purposes of calculating residual values.