

FCP-3 return Check Condition when FCP_DL too small

(T10/04-393r4)

Erich Oetting, STK

Dave Peterson, CNT

When an FCP data transfer is cut short due to incorrect setting of FCP_DL, the application client often is not informed of the error. In addition, no text exists specifying the behavior when the rddata and wrdata bits are not set correctly.

To address this problem, we propose the following changes to FCP-3.

9.1.2.7 RDDATA and WRDATA

If the RDDATA bit is set to one, the initiator expects FCP_DATA IUs for the task to be in the direction opposite to the direction of the FCP_CMND IU. This is a SCSI read operation.

If the WRDATA bit is set to one, the initiator expects FCP_DATA IUs for the task to be in the same direction as the FCP_CMND IU. This is a SCSI write operation.

If the RDDATA bit and WRDATA bit are both set to one, the initiator expects both a SCSI read operation and a SCSI write operation. This is a bidirectional SCSI ~~operation~~ command. The FCP_BIDIRECTIONAL_READ_DL field shall be included in the FCP_CMND IU payload. The initiator shall not set both the RDDATA bit and the WRDATA bit to one except for a bidirectional SCSI command.

If the RDDATA bit and WRDATA bit are both set to zero, there shall be no FCP_DATA IUs and the FCP_DL field shall be set to zero.

The target shall terminate a command and set the RSP_CODE field to "FCP_CMND fields invalid" if the following protocol errors are detected:

- a) a command requested a read operation and the RDDATA bit is set to zero or the WRDATA bit is set to one;
- b) a command requested a write operation and the WRDATA bit is set to zero or the RDDATA is set to one; or
- c) the RDDATA or WRDATA bits are set to zero for a bidirectional SCSI command.

9.3.2 FCP_DATA IUs for SCSI read and SCSI write operations

During any data transfer, the initiator shall have available a buffer of length FCP_DL. The buffer contains data to be transferred to the target if the operation is a write operation (i.e., an operation that uses the Data Out action, IU T6). The buffer receives the data if the operation is a read operation (i.e., an operation that uses the Data In action, IU I3). The target shall never request or deliver data outside the buffer length defined by FCP_DL. If the command requested that data beyond FCP_DL be transferred, the FC_RSP IU shall contain the FCP_RESID_OVER bit (see 9.4.7) set to one; and the target shall:

- a) process the command normally except that data beyond the FCP_DL count shall not be transferred;
- b) return CHECK CONDITION with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT; or
- ~~c) return an FCP_RSP_CODE of "FCP_CMND fields invalid".~~

Editor's Note - 1: Fail to see the reasoning for completing the command normally especially given the text in 9.4.11 that states "If a condition of FCP_RESID_OVER is detected, the termination state of the FCP I/O operation is not certain. Data may

or may not have been transferred and the SCSI status byte may or may not provide correct command completion information”.

During a write operation that is not using FCP_XFER_RDY IUs, the initiator indicates that it has transferred all the required data by transferring initiative to the target. The initiator shall not transfer data outside the buffer length defined by FCP_DL. If the write operation requires a total amount of data less than the amount of data provided by the initiator, the target shall discard the excess bytes. Because there were fewer bytes provided than required by FCP_DL, the FCP_RESID_UNDER bit shall be set to 1 in the FCP_RSP IU. The command is completed according to the rules specified by the SCSI command set for that command.

If the amount of data returned does not match the number of bytes of data calculated from the value of FCP_DL and the value of FCP_RESID (see 9.4.11), 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 or other failure indication. The mechanism an initiator 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.

9.3.3 FCP_DATA IUs for bidirectional SCSI ~~operations~~ commands

During any write data transfer for a bidirectional SCSI command (i.e., a command that uses the Data Out action, IU T6), the initiator always has available a buffer of the length specified by the FCP_DL field containing data to be transferred to the target.

During any read data transfer for a bidirectional SCSI command (i.e., a command that uses the Data In action, IU I3), the initiator always has available a buffer of the length specified by the FCP_BIDIRECTIONAL_READ_DL field that receives the data.

The target shall never request or deliver data outside the buffer length defined by FCP_DL or FCP_BIDIRECTIONAL_READ_DL.

If a command requests that data beyond FCP_DL be transferred, the FCP_RESID_OVER bit shall be set to one in the FCP_RSP IU: ~~and the target shall: The command is completed normally except that data beyond the count in the FCP_DL field is not transferred and the appropriate overrun condition is presented (see 9.4.7).~~

- a) process the command normally except that data beyond the FCP_DL count shall not be transferred;
- b) return CHECK CONDITION with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT; or
- ~~c) return an FCP_RSP_CODE of “FCP_CMND fields invalid”.~~

If a command requests that data beyond the value specified in the FCP_BIDIRECTIONAL_READ_DL field be transferred, the FCP_BIDIRECTIONAL_READ_RESID_OVER bit shall be set to one in the FCP_RSP IU: ~~and the target shall: The command is completed normally except that data beyond the FCP_BIDIRECTIONAL_DL_COUNT is not transferred and the appropriate overrun condition is presented (see 9.4.4).~~

- a) process the command normally except that data beyond the FCP_BIDIRECTIONAL_READ_DL count shall not be transferred;
- b) return CHECK CONDITION with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT; or
- ~~c) return an FCP_RSP_CODE of “FCP_CMND fields invalid”.~~

If the amount of data transferred does not match the value specified by the FCP_DL field as modified by any residual values recorded in the FCP_RESID field for the write data transfer of a bidirectional SCSI command or the value specified by the FCP_BIDIRECTIONAL_READ_DL field as modified by any residual values recorded in the FCP_BIDIRECTIONAL_READ_RESID field

for the read data transfer of a bidirectional SCSI command, the FCP I/O operation may be terminated with a recovery abort or other failure indication.

9.4.9 FCP_RSP_LEN_VALID

An FCP_RSP_LEN_VALID bit of one indicates that the FCP_RSP_INFO field contains valid information and that the FCP_RSP_LEN field is valid and non-zero and that it contains the count of bytes in the FCP_RSP_INFO field. The application client should examine the FCP_RSP_INFO field to determine whether or not an error condition occurred. **When the FCP_RSP_LEN_VALID bit is one, the content of the SCSI STATUS CODE field is not reliable and shall be ignored by the initiator.**

For task management functions transmitted to the target using an FCP_CMND IU, the FCP_RSP_LEN_VALID bit shall be set to one, the FCP_RSP_LEN field shall be set to the specified value, and the information in the RSP_CODE field shall indicate the completion status of the task management function.

An FCP_RSP_LEN_VALID bit of zero indicates that the FCP_RSP_LEN field is not valid and shall be treated as if its value were zero. When the FCP_RSP_LEN_VALID bit is zero, the FCP_RSP_INFO field shall have a length of zero and shall not be present.

9.4.11 FCP_RESID

If the FCP_RESID_UNDER bit or the FCP_RESID_OVER bit is set to one, the FCP_RESID field contains a count of the number of residual data bytes that were not transferred in the FCP_DATA IUs for this SCSI command. For bidirectional SCSI commands, the FCP_RESID field contains a count of the number of residual data bytes that were not transferred in the Data Out FCP_DATA IUs. For bidirectional SCSI commands, the FCP_BIDIRECTIONAL_READ_RESID field contains the corresponding count for Data In FCP_DATA IUs.

Upon successful completion of an FCP I/O operation, the residual value is normally zero and the FCP_RESID value is not valid. FCP devices having indeterminate data lengths may have a nonzero residual byte count after completing valid operations. **Targets are not required to verify that the data length implied by the contents of the CDB cause an overrun or underrun before beginning execution of a SCSI command.**

If the FCP_RESID_UNDER bit is set to one, a transfer that did not fill the buffer to the expected displacement FCP_DL was performed and the value of FCP_RESID is defined as follows:

$$\text{FCP_RESID} = \text{FCP_DL} - (\text{highest offset of any byte transmitted} + 1)$$

A condition of FCP_RESID_UNDER may not be an error for some FCP devices and some commands.

If the FCP_RESID_OVER bit is set to one, the transfer was truncated because the data transfer required by the SCSI command extended beyond the displacement value of FCP_DL. Those bytes that could be transferred without violating the FCP_DL value may be transferred. The FCP_RESID is defined as follows:

$$\text{FCP_RESID} = (\text{transfer length required by command}) - \text{FCP_DL}$$

If a condition of FCP_RESID_OVER is detected, the termination state of the FCP I/O operation is not certain. Data may or may not have been transferred and the SCSI status byte may or may not provide correct command completion information.

If the FCP_RESID_UNDER and the FCP_RESID_OVER bits are set to zero, the FCP_RESID field is not meaningful and may have any value. The FCP_RESID field is always included in the FCP_RSP IU.

NOTE 5 – Some early target implementations presented the FCP_RSP IU without the FCP_RESID, FCP_SNS_LEN, and FCP_RSP_LEN fields if the FCP_RESID_UNDER, FCP_RESID_OVER, FCP_SNS_LEN_VALID, and FCP_RSP_LEN_VALID bits were all set to zero. Initiators should be tolerant of this non-standard behavior.

9.4.12 FCP_BIDIRECTIONAL_READ_RESID

The FCP_BIDIRECTIONAL_READ_RESID field is included in the FCP_RSP IU for all bidirectional SCSI commands. If the FCP_BIDI_READ_RESID_UNDER bit or the FCP_BIDI_READ_RESID_OVER bit is set to one, the FCP_BIDIRECTIONAL_READ_RESID

field contains a count of the number of residual data bytes that were not transferred in the read FCP_DATA IUs for this bidirectional SCSI command.

Upon successful completion of an FCP I/O operation, the residual value is normally zero and the FCP_BIDIRECTIONAL_READ_RESID value is not valid. FCP devices having indeterminate data lengths may have a nonzero residual byte count after completing valid operations. **Targets are not required to verify that the data length implied by the contents of the CDB cause an overrun or underrun before beginning execution of a SCSI command.**

If the FCP_BIDI_READ_RESID_UNDER bit is set to one, a transfer that did not fill the buffer to the expected displacement FCP_BIDIRECTIONAL_READ_DL was performed and the value of FCP_BIDIRECTIONAL_READ_RESID is defined as follows:

$$\text{FCP_BIDIRECTIONAL_READ_RESID} = \text{FCP_BIDIRECTIONAL_READ_DL} - (\text{highest offset of any byte transmitted} + 1)$$

A condition of FCP_BIDI_READ_RESID_UNDER may not be an error for some FCP devices and some commands.

If the FCP_BIDI_READ_RESID_OVER bit is set to one, the transfer was truncated because the data transfer required by the SCSI command extended beyond the displacement value of FCP_BIDIRECTIONAL_READ_DL. Those bytes that could be transferred without violating the FCP_DL value may be transferred. The FCP_BIDIRECTIONAL_READ_RESID is defined as follows:

$$\text{FCP_BIDIRECTIONAL_READ_RESID} = (\text{read transfer length required by command}) - \text{FCP_BIDIRECTIONAL_READ_DL}$$

If a condition of FCP_BIDI_READ_RESID_OVER is detected, the termination state of the FCP I/O operation is not certain. Data may or may not have been transferred and the SCSI status byte may or may not provide correct command completion information.

If the FCP_BIDI_READ_RESID_UNDER and the FCP_RESID_OVER bits are both set to zero, the FCP_BIDIRECTIONAL_READ_RESID field is not meaningful and may have any value.

4.3 Bidirectional and unidirectional operation

A device server that supports bidirectional operation may implement some commands as unidirectional commands as well as some commands that provide for bidirectional transfer. Two FCP_RSP IU formats are defined. For commands that specify bidirectional transfer by setting both the RDDATA and WRDATA bits to one, the bidirectional FCP_RSP IU payload shall be used for presenting all status and error conditions. For commands that specify either no transfer or unidirectional transfer by setting either the RDDATA or WRDATA bit or both to zero, the unidirectional FCP_RSP IU payload shall be used for presenting all status and error conditions. The format of the FCP_RSP IU that is returned depends only on the state of the RDDATA and WRDATA bits and is not influenced by whether the command itself requests bidirectional behavior.

A device server that does not support bidirectional operation shall use the unidirectional FCP_RSP IU payload for presenting all status and error conditions. If a device server that does not support bidirectional operation receives a command that requests bidirectional operation by setting both the RDDATA and WRDATA bits to one, the device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of **ILLEGAL FIELD IN PARAMETER LIST INVALID FIELD IN COMMAND INFORMATION UNIT**.

