

FCP-3 return Check Condition when FCP_DL too small

(T10/05-044r3)

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. Further, tape drives typically do not want to transfer any data if there is a data length mismatch while disk drives typically want to continue to process the command. In addition, no text exists specifying the behavior when the RDDATA and WRDATA bits are not set correctly.

To address this problem, the following changes to FCP-3 are proposed:

9.1.2.7 RDDATA and WRDATA

If the RDDATA bit is set to one, the initiator expects to receive FCP_DATA IUs from the target. This is a SCSI read operation.

If the WRDATA bit is set to one, the initiator expects to transmit FCP_DATA IUs to the target. 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 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 return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT if the following protocol errors are detected:

- a) a read operation has the RDDATA bit set to zero or the WRDATA bit set to one;
- b) a write operation has the WRDATA bit set to zero or the RDDATA bit set to one; or
- c) a bidirectional SCSI command has either the RDDATA bit set to zero or the WRDATA bit set to zero.

NOTE – Targets compliant to previous versions of this standard may terminate the command and return an FCP_RSP IU with the RSP_CODE field to “FCP_CMND fields invalid” (see table 24) for some protocol errors.

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 not request or deliver data outside the buffer length defined by FCP_DL.

If the command requested that data beyond FCP_DL be transferred, the target shall set the FCP_RESID_OVER bit (see 9.4.8) to one in the FCP_RSP IU and shall:

- a) process the command normally except that data beyond the FCP_DL count shall not be requested or transferred; or
- b) transfer no data and return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

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 Sequence 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 (see 9.4.7) shall be set to **one** 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 **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 9.4.12), 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 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 commands

During a bidirectional command, the initiator shall always have available a buffer with the length specified by the FCP_DL field to transfer data to the target. The target shall not request data outside the buffer length specified by the FCP_DL field.

During a bidirectional command, the initiator shall always have available a buffer with the length specified by the FCP_BIDIRECTIONAL_DL field to receive data from the target. The target shall not transfer data outside the buffer length specified by the FCP_BIDIRECTIONAL_DL field.

If a **bidirectional** command requested that data beyond FCP_DL be transferred, the target shall set the FCP_RESID_OVER bit (see 9.4.8) to one in the FCP_RSP IU and shall:

- a) process the command normally except that data beyond the FCP_DL count shall not be transferred; or
- b) transfer no data in either direction and return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

If a **bidirectional** command requests that data beyond the value specified in the FCP_BIDIRECTIONAL_READ_DL field be transferred, the target shall set the FCP_BIDI_READ_RESID_OVER bit (see 9.4.5) to one in the FCP_RSP IU and shall:

- a) process the command normally except that data beyond the FCP_BIDIRECTIONAL_READ_DL count shall not be transferred; or
- b) transfer no data in either direction and return CHECK CONDITION with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

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 9.4.12) for the write data transfer of a bidirectional SCSI command or the number of bytes of data calculated from the value of the FCP_BIDIRECTIONAL_READ_DL field and the value of the FCP_BIDIRECTIONAL_READ_RESID field (see 9.4.13) for the read data transfer of a bidirectional SCSI command, the FCP I/O operation may be terminated with a recovery abort (see 12.3) 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.

9.4.9 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

For read commands and write operations, if the FCP_RESID_UNDER bit is set to one, the FCP_RESID field contains a count of the number of residual data bytes that were not transferred in FCP_DATA IUs.

For read commands and write operations, if the FCP_RESID_OVER bit is set to one, the FCP_RESID field contains the excess of the number of bytes required by the SCSI command to be transferred over the number of bytes specified by the FCP_DL field.

For bidirectional SCSI commands, if the FCP_RESID_UNDER 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 Data-Out FCP_DATA IUs for the command.

For bidirectional SCSI commands, if the FCP_RESID_OVER bit is set to one, the FCP_RESID field contains the excess of the number of bytes required to be transferred in the Data-Out FCP_DATA IUs by the command over the number of bytes specified in the FCP_DL field.

For bidirectional SCSI commands, the FCP_BIDIRECTIONAL_READ_RESID field (see 9.4.13) 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, refer to 9.3.2 or 9.3.3. The FCP_RESID value is defined as follows:

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

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.

For bidirectional commands, if the FCP_BIDI_READ_RESID_UNDER 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 Data-In FCP_DATA IUs.

For bidirectional commands, if the FCP_BIDI_READ_RESID_OVER bit is set to one, the FCP_BIDIRECTIONAL_READ_RESID field contains the excess of the number of bytes required by the command to be transferred in Data-In FCP_DATA IUs over the number of bytes specified by the FCP_BIDIRECTIONAL_READ_DL field.

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, refer to 9.3.3. The FCP_BIDIRECTIONAL_READ_RESID value is defined as follows:

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

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 commands and FCP_RSP IU format

A device server that supports bidirectional commands may implement both unidirectional and bidirectional commands. Two FCP_RSP IU formats are defined. For commands that set 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 set 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.

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