

FCP-3: FCP_DATA IU Clarification

(T10/04-387r0)

Dave Peterson, CNT

There are several instances in the FCP standard where it implies that only a single FCP_DATA is sent in response to an FCP_XFER_RDY. This is clearly not correct since multiple FCP_DATA IUs may be sent in response to an FCP_XFER_RDY.

4.2 Device management

.
. .
.

When the device server for the command has completed the interpretation of the command, has determined that write data transfer is required, and is prepared to request the data delivery service, it sends a data descriptor IU containing the FCP_XFER_RDY IU payload to the initiator to indicate which portion of the data is to be transferred. The FCP_Port that is the initiator then transmits a solicited data IU to the target containing the FCP_DATA IU payload requested by the FCP_XFER_RDY IU. The FCP_XFER_RDY IU and FCP_DATA IU payloads constitute the Receive Data-Out protocol service request and Data-Out Received service confirmation described in SAM-3. Data delivery requests containing FCP_XFER_RDY IU and returning FCP_DATA IU payloads continue until the data transfer requested by the SCSI command is complete. **One FCP_DATA IU shall follow each FCP_XFER_RDY IU.** If the system has mechanisms outside the scope of this standard for controlling the data transfer length, the transmission of the initial FCP_XFER_RDY IU may be disabled (see 10.2.10).

9.2.1 Overview and format of FCP_XFER_RDY IU

The FCP_XFER_RDY IU indicates that the target is prepared to receive part or all of the data for a write command. The FCP_XFER_RDY IU contains those parameters of the SAM-3 data delivery service required by the initiator, including the length and beginning relative offset **of the FCP_DATA IU that is requested.** Since the target has established buffering and caching resources based on the requested data, the initiator shall provide the described data **in the requested FCP_DATA IU.** The initiator shall be ready to transmit any part or all of the number of bytes indicated in the FCP_DL field if requested.

FCP_XFER_RDY IUs shall be transmitted preceding each write FCP_DATA IU when the WRITE FCP_XFER_RDY DISABLED bit is set to zero by PRLI. If the target and initiator have negotiated write FCP_XFER_RDY disabled, FCP_XFER_RDY IUs shall be transmitted to request each write FCP_DATA IU after the first FCP_DATA IU of the command. The first FCP_DATA IU is transmitted without a preceding FCP_XFER_RDY IU (see 6.3.4).

9.2.3 FCP_BURST_LEN

The FCP_BURST_LEN field contains a value indicating the amount of buffer space prepared for all bytes to be transferred **in the next FCP_DATA IU and requests the transfer from the initiator of an IU of that length.** The value in the FCP_BURST_LEN field is the same as the SCSI data delivery request byte count (see SAM-3).

The value in the FCP_BURST_LEN field shall not exceed the maximum burst length defined by the disconnect-reconnect page of MODE SELECT and MODE SENSE commands. See 10.2.7. The sum of the value of FCP_BURST_LEN field and the value of FCP_DATA_RO shall not exceed the value of FCP_DL. The value in the FCP_BURST_LEN field shall not be zero.

9.3.1 FCP_DATA IU Overview

The data associated with a particular FCP I/O operation is transmitted in the same exchange that sent the FCP_CMND IU requesting the transfer.

SCSI data transfers may be performed by one or more data delivery requests, each one performing a transfer no longer than the maximum burst length defined by the parameters of the disconnect-reconnect page of the MODE SENSE and MODE SELECT commands.

If more than one FCP_DATA IU is used to transfer the data, the relative offset value in the PARAMETER field is used to ensure that the SCSI data is reassembled in the proper order (see 5.4.2.12). If an FCP_XFER_RDY IU is used to describe a data transfer and the first frame of the requested FCP_DATA IU has a relative offset that differs from the value in the FCP_DATA_RO field of the FCP_XFER_RDY IU, the target shall post the error code "FCP_DATA Parameter mismatch with FCP_DATA_RO" in the FCP_RSP_INFO field of the FCP_RSP IU.

If required by the PRLI FCP service parameters, each Data Out action FCP_DATA IU shall be preceded by an FCP_XFER_RDY IU containing a standard data descriptor payload that indicates the location and length of the data delivery. If the PRLI FCP Service Parameters specify WRITE FCP_XFER_RDY DISABLED, the first FCP_DATA IU shall be transmitted without a preceding FCP_XFER_RDY IU.

Table 6 - FCP Information Units (IUs) sent to targets

IU	SCSI-3 primitive	Data block		F/M/L	SI	M/O
		CAT	Content			
T1	Command / Task Mgmt Rqst	6	FCP_CMND	F	T	M
T2	Command request	6	FCP_CMND	F	H	O
T3	Command request (Linked)	6	FCP_CMND	M	T	O
T4	Command request (Linked)	6	FCP_CMND	M	H	O
T6	Data Out action	1	FCP_DATA	M	T	M
T12	Confirm	3	none	L	T	O

Notes:
T5, T7, T8, T9, T10, and T11 are obsolete.
T2 and T4 are only permitted when transfer ready IUs are disabled (see table 9).
T3 and T4 are only permitted for linked SCSI commands.
T2 and T4 allow optional sequence streaming during write operations.
T12 is only permitted in response to an I5 frame requesting the confirmed completion protocol.
See table 7

Key:

IU	Information Unit identifier
CAT	Information category of Device_Data frames carrying the data block (see FC-FS-2)
CONTENT	Contents (payload) of data block
F/M/L	First/Middle/Last Sequence of Exchange (FC-FS-2)
	F First
	M Middle
	L Last
SI	Sequence Initiative: Held or Transferred (FC-FS-2)
	H Held
	T Transferred
M/O	Mandatory/Optional Sequence
	M Mandatory
	O Optional

Word 3, Bit 0: WRITE FCP_XFER_RDY DISABLED: When the WRITE FCP_XFER_RDY DISABLED bit is set to zero, FCP_XFER_RDY IUs shall be transmitted by the target to request each of the SCSI write FCP_DATA IUs from the initiator. When the WRITE FCP_XFER_RDY DISABLED bit is set to one, FCP_XFER_RDY IUs shall not be used before the first FCP_DATA IU to be transferred in the SCSI write operation. If both the Originator and Responder choose to disable write FCP_XFER_RDY IUs, then all FCP I/O operations performing SCSI writes between the FCP_Ports shall operate without using the FCP_XFER_RDY IU before the first FCP_DATA IU. The FCP_XFER_RDY IU shall be transmitted to request each additional FCP_DATA IU, if any. If either the Originator or the Responder requires the use of FCP_XFER_RDY IUs during SCSI writes, then the Exchange Responder shall transmit an FCP_XFER_RDY IU requesting each FCP_DATA IU, including the first, from the Exchange Originator.

10.2.10 FIRST BURST SIZE

When write transfer ready is disabled, the FIRST BURST SIZE field indicates the maximum amount of all bytes that shall be transmitted in the first FCP_DATA IU sent from the initiator to the target. If all data is transmitted in the first IU, no subsequent FCP_XFER_RDY IUs shall be transmitted by the target. If the maximum amount of data has been transmitted, but more data remains to be transferred, the target shall request that data with subsequent FCP_XFER_RDY IUs.

When write transfer ready is enabled, the FIRST BURST SIZE field is ignored and permission to transmit data from the initiator to the target is managed using FCP_XFER_RDY IUs. For data transmissions from the target to the initiator, the FIRST BURST SIZE field is ignored.

The FIRST BURST SIZE field value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes). A value of zero indicates that there is no first burst size limit. The FIRST BURST SIZE field shall be implemented by all FCP devices that support the disabling of write transfer ready. The application client and device server may use the value of this parameter to adjust internal maximum buffering requirements.

10.2.7 MAXIMUM BURST SIZE FIELD

The MAXIMUM BURST SIZE field indicates the maximum size of all bytes in an FCP_DATA IU that the device server shall transfer to the initiator or request from the initiator. This parameter does not affect how much data is transferred in a single interconnect tenancy. This value is expressed in increments of 512 bytes (e.g., a value of 1 means 512 bytes, two means 1024 bytes, etc.). The device server may round this value down as defined in SPC-3. A value of zero indicates there is no limit on the amount of data transferred per data transfer operation. This value shall be implemented by all FCP devices. The application client and device server may use the value of this parameter to adjust internal maximum buffering requirements.