

**To:** T10 Technical Committee  
**From:** Rob Elliott, Compaq Computer Corporation (Robert.Elliott@compaq.com)  
 Carl Zeitler, Compaq Computer Corporation (Carl.Zeitler@compaq.com)  
**Date:** 26 September 2000  
**Subject:** Bidirectional data transfers in FCP-2

T10/00-309r1 by George Penokie, accepted into SAM-2 revision 14, modified SAM to allow protocols and commands to support bidirectional data transfers – commands which transfer both read and write data. This proposal modifies FCP-2 (revision 4a) to support such transfers.

Although Fibre Channel is a full duplex interconnect, data transfers within an exchange (command) are half-duplex due to the sequence initiative concept. This means only a write or read for the command can be active at one time, not both. This is an FC-FS issue, not an FCP-2 issue. This proposal does not try to eliminate this restriction and limits transfers to half-duplex.

The target has full control over when read or write data is transferred. It can run them in any order, and switch as many times as it wants. If it wants to switch from a write to a read, it just sends a FCP\_DATA IU. If it wants to switch from a read to a write, it sends an FCP\_XFER\_RDY IU with Sequence Initiative transfer indicated. The initiator follows by sending an FCP\_DATA IUs, with Sequence Initiative transfer indicated. No additional IUs are needed.

The Command IU changes to allow both READDATA and WRITEDATA bits to be enabled at once. An additional Data Length (DL) field is added when both READDATA and WRITEDATA are set since the read data and write data sizes may be different. The data structure is already variable length thanks to the ADDITIONAL FCP CDB LENGTH field, so 4 more bytes should be tolerable.

The ability to disable the initial FCP\_XFER\_RDY via process login creates a problem for a bidirectional command which needs to transfer read data first. If the initiator is allowed to immediately send write data after sending the command, the target would need to buffer it. To avoid this issue, the existing WRITE FCP\_XFER\_RDY DISABLE bit is only applied to unidirectional write commands and a new bit is added for bidirectional commands. The target may enable the bit if it doesn't support any bidirectional commands that require read data first, or if it has adequate buffer capability to handle the issue. (Alternatively, disabling FCP\_XFER\_RDY for bidirectional commands could simply be disallowed. Is this feature used by real devices?)

Overrun and underrun of the Data Length (FCP\_DL) complicates the RSP IU. The read and write could each result in an error. The existing bits cover the write direction; extra bits are added to indicate these errors in the read direction. An extra field is added for the bidirectional read residual count if either of the bidirection read error bits is set. This makes the field size vary based on presence of an error. Alternatively, it could be fixed at a larger size for bidirectional commands. Alternatively, this type of error reporting could simply not be supported for bidirectional commands (is this feature used by real devices?)

“BIDIRECTIONAL\_READ” was added to the names of existing fields when creating new ones. A shorter acronym may be preferred (maybe just “BIDI”).

#### **Section 6.3.7.1 FCP Process Login request service parameter page format**

The FCP service parameter page for the Process Login request is shown in table 12.

Add an entry:

<b>FCP service parameter</b>	<b>Word</b>	<b>Bit</b>
Bidirectional Read	3	10
XFER_RDY Disabled		

**Section 6.3.7.15 Word 3, Bit 0: WRITE XFER\_RDY DISABLED**

When this bit is set to 0, FCP\_XFER\_RDY IUs shall be used for SCSI write operations *for unidirectional write commands*. When this bit is set to 1, FCP\_XFER\_RDY IUs may ~~be not~~ be used before the first FCP\_DATA IU to be transferred in the 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, after the first one. If either the originator or the responder requires the use of FCP\_XFER\_RDY IUs during writes, then the exchange responder shall transmit an FCP\_XFER\_RDY IU requesting each FCP\_DATA IU, including the first, from the exchange originator.

***[new] Section 6.3.7.xx Word 3, Bit 10: BIDIRECTIONAL WRITE XFER\_RDY DISABLED***

*When this bit is set to 0, FCP\_XFER\_RDY IUs shall be used for write operations for SCSI bidirectional commands.*

*If both the originator and responder set this bit to 1, then all SCSI bidirectional commands performing SCSI write operations 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 write FCP\_DATA IU, if any, after the first one.*

*If either the originator or the responder requires the use of FCP\_XFER\_RDY IUs during writes, then the exchange responder shall transmit an FCP\_XFER\_RDY IU requesting each FCP\_DATA IU, including the first, from the exchange originator.*

**Section 9.1 FCP\_CMND IU**

Add a FCP BIDIRECTIONAL READ DL field to the table:

**Table 24 – FCP\_CMND payload**

0	(MSB)	FCP_LUN			(LSB)
7					
8		COMMAND REFERENCE NUMBER			
9		RESERVED	TASK ATTRIBUTE		
10		TASK MANAGEMENT FLAGS			
11		RESERVED	ADDITIONAL FCP CDB LENGTH	RDDATA	WRDATA
12	(MSB)	FCP CDB			(LSB)
27					
28	(MSB)	ADDITIONAL FCP CDB			(LSB)
n+1	(MSB)	FCP DL			(LSB)
n+4					
n+5	(MSB)	<i>FCP BIDIRECTIONAL READ DL</i>			
n+8					(LSB)

**Section 9.1.1.4 Task management flags, Byte 10**

Task management function shall be transmitted by the initiator (Exchange Originator) using a new Exchange. If any task management flag is set to 1, the FCP\_CDB, FCP\_DL, *FCP\_BIDIRECTIONAL\_READ\_DL*, TASK ATTRIBUTES, RDDATA, and WRDATA fields and bits are not valid and are ignored. No more than one task management flag shall be set to 1 in any FCP\_CMND IU.

#### **Section 9.1.1.6 Read Data, Byte 11**

*[typo note: 9.1.1.x headers use a mix of small caps and mixed case]*

**READ DATA**, when set to 1, specifies that 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 ~~type or~~ *bidirectional* operation.

#### **Section 9.1.1.7 WRITE DATA, BYTE 11**

**WRITE DATA**, when set to 1, specifies that 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 *or bidirectional* operation. If both READ DATA and WRITE DATA are set to 0, there shall be no FCP\_DATA IUs and FCP\_DL shall be 0. ~~The initiator shall not set both the READ DATA and the WRITE DATA bits to 1. If both READ DATA and WRITE DATA are set to 1, the command is a bidirectional command and the FCP\_BIDIRECTIONAL\_READ\_DL field shall be included in the FCP\_CMND payload.~~

#### **Section 9.1.1.10 FCP\_DL**

*For a bidirectional command, the FCP\_DL field contains a count of the greatest number of data bytes expected to be transferred from the application client data buffer by the SCSI command. The parameter is the command byte count defined by SAM-2.*

*For a unidirectional write command, ~~The~~ the FCP\_DL field contains a count of the greatest number of data bytes expected to be transferred ~~to or~~ from the application client data buffer by the SCSI ~~CDB~~ command. For a unidirectional read command, the FCP\_DL field contains a count of the greatest number of data bytes expected to be transferred to the application client data buffer by the SCSI command. The parameter is the command byte count defined by SAM-2. [break paragraph here]*

An FCP\_DL value of 0 indicates that no data transfer is expected regardless of the state of the READ DATA and WRITE DATA bits and that no FCP\_XFER\_RDY or FCP\_DATA IUs shall be transferred.

#### **[new] Section 9.1.1.11 FCP\_BIDIRECTIONAL\_READ\_DL**

*If RDDATA and WRDATA are both set to 1, a FCP\_BIDIRECTIONAL\_READ\_DL field follows the FCP\_DL field. The FCP\_BIDIRECTIONAL\_READ\_DL field contains a count of the greatest number of data bytes expected to be transferred to the application client data buffer by the SCSI command. The parameter is the command byte count defined by SAM-2. An FCP\_BIDIRECTIONAL\_READ\_DL value of 0 indicates that no read data transfer is expected regardless of the state of the READ DATA bit and that no FCP\_DATA IUs shall be transferred for read data.*

*If either RDDATA or WRDATA is set to 0, the FCP\_BIDIRECTIONAL\_READ\_DL field shall not be included in the FCP\_CMND\_IU data payload.*

#### **Section 9.3 FCP\_DATA IU**

...  
During any *write* data transfer (*an operation that uses Data Out actions, IUs T6 or T7*), the initiator shall always have available a buffer of length FCP\_DL. ~~The buffer contains~~ *containing* data to be transferred to the target ~~if the operation is a write operation (an operation that uses Data Out actions, IUs T6 or T7).~~

During any read data transfer for a unidirectional read command (an operation that uses the Data In action, IU I3), the initiator shall always have available a buffer of length FCP\_DL that receives the data if the operation is a read operation (an operation that uses the Data In action, IU I4).

[note: existing text above refers to I4 FCP\_RSP when it should refer to I3 FCP\_DATA]

During any read data transfer for a bidirectional command (an operation that uses the Data In action, IU I3), the initiator shall always have available a buffer of length FCP\_BIDIRECTIONAL\_READ\_DL 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 the command requested that data beyond FCP\_DL be transferred, the FC\_RSP IU shall contain the FCP\_RESID\_UNDER bit. If the command requested that data beyond FCP\_BIDIRECTIONAL\_READ\_DL be transferred, the FC\_RSP IU shall contain the FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER bit. The command is completed normally except for presentation of the overrun condition. See 9.4.23. [this changed going from fcp2r4 to fcp2r4a – it should point to the FCP\_RESID\_UNDER bit]

...  
 If the amount of data returned [? Does this mean the paragraph applies to read data only ?] does not match FCP\_DL for a unidirectional read command or FCP\_BIDIRECTIONAL\_READ\_DL for a bidirectional command, 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 a SCSI Initiator uses to determine that the correct amount of data has been returned is outside the scope of this standard. Data that has been retransmitted and overlaid shall be counted only once.

**Section 9.4.1 Overview and format of FCP\_RSP IU**

Add two bits and a field to Table 28 FCP\_RSP payload:

**Table 24 – FCP\_RSP payload**

0-9	Reserved							
10	R s v d	FCP BIDIRECT IONAL READ RESID UNDER	FCP BIDIRECT IONAL READ RESID OVER	FCP CONF REQ	FCP RESID UNDER	FCP RESID OVER	FCP SNS_LEN VALID	FCP RSP_LEN VALID
11		SCSI status code						
12-15	FCP_RESID							
16-19	FCP_SNS_LEN (= n)							
20-23	FCP_RSP_LEN (= m)							
24	FCP_RSP_INFO (m bytes long)							
23+m								
24+m	FCP_SNS_INFO (n bytes long)							
23+m+n								
24+m+n	FCP_BIDIRECTIONAL_READ_RESID							
27+m+n								

**[new] Section 9.4.x FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER**

FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER, when 1, indicates that the FCP\_BIDIRECTIONAL\_READ\_RESID field is present and valid and contains the count of bytes that were expected to be transferred, but were not transferred. The application client should examine the FCP\_BIDIRECTIONAL\_READ\_RESID field in the context of the command to determine whether or not an error condition occurred.

**[new] Section 9.4.x FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER**

*FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER, when 1, indicates that the FCP\_BIDIRECTIONAL\_READ\_RESID field is present and valid and contains the count of bytes that could not be transferred because the FCP\_BIDIRECTIONAL\_READ\_DL was not sufficient. The application client should examine the FCP\_RESID field in the context of the command to determine whether or not an error condition occurred.*

**[no change] Section 9.4.3 FCP\_RESID\_UNDER**

*FCP\_RESID\_UNDER, when 1, indicates that the FCP\_RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred. The application client should examine the FCP\_RESID field in the context of the command to determine whether or not an error condition occurred.*

**[no change] Section 9.4.4 FCP\_RESID\_OVER**

*FCP\_RESID\_OVER, when 1, indicates that the FCP\_RESID field is valid and contains the count of bytes that could not be transferred because the FCP\_DL was not sufficient. The application client should examine the FCP\_RESID field in the context of the command to determine whether or not an error condition occurred.*

**[new] 9.4.x FCP\_BIDIRECTIONAL\_READ\_RESID**

*If either the FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER bit or the FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER bit is 1, the FCP\_BIDIRECTIONAL\_READ\_RESID field shall be included in the FCP\_RSP payload and shall contain a count of the number of residual data bytes that were not transferred in the FCP\_DATA IUs for this bidirectional SCSI command. Upon successful completion of a FCP I/O operation, the residual value is normally 0 and the FCP\_BIDIRECTIONAL\_READ\_RESID value is not valid. 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 an SCSI command.*

*If the FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER bit is set, 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 a number equal to:*

*FCP\_BIDIRECTIONAL\_READ\_DL - highest offset of any byte transmitted*

*A condition of FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER may not be an error for some devices and some commands.*

*If the FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER bit is set, 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 a number equal to:*

*(Transfer length required by command) - FCP\_BIDIRECTIONAL\_READ\_DL*

*If a condition of FCP\_BIDIRECTIONAL\_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\_BIDIRECTIONAL\_READ\_RESID\_UNDER and the FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER bits are 0, the FCP\_BIDIRECTIONAL\_READ\_RESID field is not meaningful and may have any value.*

*The FCP\_BIDIRECTIONAL\_READ\_RESID field is only included in the FCP\_RSP IU if either FCP\_BIDIRECTIONAL\_READ\_RESID\_OVER or FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER is set.*

### **Section 12.2.2 FCP Error Detection using protocol errors for all classes of service.**

...

d) a *unidirectional* read command completed with the data count smaller than FCP\_DL and FCP\_RESID\_UNDER ~~is~~ set to 0, *or a bidirectional read command completed with the data count smaller than FCP\_BIDIRECTIONAL\_READ\_DL and FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER set to 0;*

e) a *unidirectional* read ~~type~~ command completed with the data count smaller than FCP\_DL, FCP\_RESID\_UNDER ~~is~~ set to 1, and the data count plus FCP\_RESID ~~is~~ not equal to FCP\_DL; *or a bidirectional command completed with the read data count smaller than FCP\_BIDIRECTIONAL\_READ\_DL, FCP\_BIDIRECTIONAL\_READ\_RESID\_UNDER set to 1, and the data count plus FCP\_BIDIRECTIONAL\_READ\_RESID not equal to FCP\_BIDIRECTIONAL\_READ\_DL;* and

### **Section 12.4.1.5 FCP RSP IU Recovery**

...

An Exchange carrying a command that was terminated by a CHECK CONDITION requesting FCP\_CONF prior to transferring data may have the same REC values as an Exchange carrying a command having an FCP\_XFER\_RDY IU not received by the initiator. For a write *or bidirectional* command with a non-zero FCP\_DL, the parameters for the SRR shall indicate that an FCP\_XFER\_RDY IU is expected from the target. The target is aware of the actual present state of the transfer and response and shall either retry the FCP\_XFER\_RDY IU or, if the actual data transfer length for the command was zero, retry the FCP\_RSP.

### **Section A.2 Application client SCSI command services**

The SCSI command services shall be requested by the application client using a procedure call defined as:

Service response = execute command (  
fully qualified exchange identifier + logical unit number,  
command descriptor block,  
[task attribute],  
[data-out buffer],  
[command byte count] †  
[data-in buffer],  
status).

### **Section A.4.1 Overview of data buffer movement services**

The SCSI data buffer movement services shall be requested from the device server using a procedure call defined as:

Service response = move data buffer (  
fully qualified exchange identifier + logical unit number,  
device server buffer,  
application client buffer offset,  
request byte count ||).

~~Only one type of data buffer movement procedure call shall be used while processing one command, either data-in delivery or data-out delivery. Either data-in delivery, data-out delivery, both data-in and data-out delivery, or no data delivery may be used while processing one command. If both are used (for a bidirectional command), the device server shall combine the data-in and data-out service responses into one service response.~~

### **[new] Section C.1.x SCSI FCI bidirectional operation with write before read**

A typical SCSI FCI bidirectional operation with a single data IU transferred in each direction is shown in table C.x. The example command accepts write data before returning read data.

Initiator function	IU	Target function
--------------------	----	-----------------

Command request	T1, FCP_CMND ->	
		[Prepare data out transfer buffer]
	<- I1, FCP_XFER_RDY	Data out delivery request
Data out action	T6, FCP_DATA ->	
		[Prepare data in transfer]
	<- I3, FCP_DATA	Data in action
		[Prepare response message]
	<- I4, FCP_RSP	Response
[Indicate command completion]		

**[new] Section C.1.x SCSI FCI bidirectional operation with read before write**

A typical SCSI FCI bidirectional operation with a single data IU transferred in each direction is shown in table C.x. The example command returns read data before accepting write data.

Initiator function	IU	Target function
Command request	T1, FCP_CMND ->	
		[Prepare data in transfer]
	<- I3, FCP_DATA	Data in action
		[Prepare data out transfer buffer]
	<- I1, FCP_XFER_RDY	Data out delivery request
Data out action	T6, FCP_DATA ->	
		[Prepare response message]
	<- I4, FCP_RSP	Response
[Indicate command completion]		

**[new] Section C.1.x SCSI FCI bidirectional operation with write before read and FCP\_XFER\_RDY disabled**

A SCSI FCI bidirectional operation with two write data IUs and one read data IU is shown in table C.x. The example command accepts write data before returning read data. The initial FCP\_XFER\_RDY IU has been disabled for bidirectional commands during process login.

Initiator function	IU	Target function
Command request	T1, FCP_CMND ->	
Data out action	FCP_DATA ->	First Data out
	<- I1, FCP_XFER_RDY	Second Data out delivery request
Data out action	T6, FCP_DATA ->	
	<- I1, FCP_XFER_RDY	Last Data out delivery request
Data out action	T6, FCP_DATA ->	
		[Prepare data in transfer]
	<- I3, FCP_DATA	Data in action
		[Prepare response message]
	<- I4, FCP_RSP	Response
[Indicate command completion]		

**[new] Section C.1.x SCSI FCI bidirectional operation with intermixed writes and reads**

A SCSI FCI bidirectional operation with three data IUs transferred in each direction is shown in table C.x. The example command accepts some write data before returning read data, but intermixes writes and reads thereafter.

Initiator function	IU	Target function
Command request	T1, FCP_CMND ->	
		[Prepare data out buffer]
	<- I1, FCP_XFER_RDY	First Data out delivery request
First Data out action	T6, FCP_DATA ->	
		[Prepare data in transfer]
	<- I3, FCP_DATA	First Data in action
	<- I1, FCP_XFER_RDY	Second Data out delivery request
Second Data out action	T6, FCP_DATA ->	
	<- I1, FCP_XFER_RDY	Last Data out delivery request
Third Data out action	T6, FCP_DATA ->	
	<- I3, FCP_DATA	Second Data in action
	<- I3, FCP_DATA	Last Data in action
		[Prepare response message]
	<- I4, FCP_RSP	Response
[Indicate command completion]		

### **Section E.1 Introduction**

This annex diagrams several error detection and recovery procedures for SCSI devices that are executing a Read command with multiple fixed blocks. The examples that follow are shown in Class 3 with in-order delivery for simplicity.

In the following examples the method of error detection and recovery will depend on the Initiator's capabilities:

If the Initiator can only detect that lost read data has occurred by comparing its internal transferred byte count with the Target reported transfer byte count (FCP\_DL - FCP\_RESID *for unidirectional read commands or FCP\_BIDIRECTIONAL\_READ\_DL - FCI\_BIDIRECTIONAL\_READ\_RESID for bidirectional commands*) then the Initiator will use the After Status Recovery Method. See E.2.

### **Section E.2.1 Discovery [no change; this is just an example and need not consider bidir]**

During the transfer of the read data a frame was lost or dropped. However this Initiator was not designed to detect a sequence error when it occurred. The Target reported in the FCP\_RSP that the transferred byte count (FCP\_DL - FCP\_RESID) was 36000. The Initiator compared this to its internal byte transfer count and discovered that some data had been lost.