Fibre Channel Protocol for SCSI

FCP REV 0

proposal for draft working document
American National Standard
for Information Systems
September 4, 1992

Secretariat:

Computer & Business Equipment Manufacturers Association

ABSTRACT: This proposal describes the frame format and protocol definitions required to transfer commands and data between a SCSI (Small Computer System Interface) Initiator and Target using a serial link interface operating according to the Fibre Channel protocol requirements.

NOTE: This is a proposal for a working document of X3T9.2 and X3T9.3, Task Groups of Accredited Standards Committee X3. As such, this is not a completed standard. The contents are actively being modified by the two Task Groups. This document is made available for review and comment only. For current information on the status of this document contact the individuals shown below:

POINTS OF CONTACT:

Don Tolmie (X3T9.3 Chairman)
Los Alamos National Laboratory
C-5, MS-B255
Los Alamos, NM 87545
(505) 667-5502
E-Mail: det@lanl.gov

John Lohmeyer (X3T9.2 Chairman)
NCR Corporation
1635 Aeroplaza Dr.
Colorado Springs, CO 80916
(719) 596-5795 x362
E-Mail: john.lohmeyer@wichitaks.ncr.com

I. Dal Allan (Fibre Channel Working Group Chair)
ENDL
14426 Black Walnut Court
Saratoga, CA 95070
(408) 867-6630
E-Mail:0002501752@mcimail.com

Paul Ramsay / Robert Snively (Editors)
Sun Microsystems, MTV16-04 / MTV15-46
2550 Garcia Ave
Mountain View, CA 94043-1100
(415) 336-4622 / (415) 336-5332
E-Mail: paul.ramsay@eng.sun.com
bob.snively@eng.sun.com

CONTENTS

Foreword

The Small Computer System Interface (SCSI) command set is widely used and applicable to a wide variety of device types. The transmission of SCSI command set information across Fibre Channel links allows the large body of SCSI application and driver software to be successfully used in the high performance serial IO channel environment.

This document describes the manner of using Fibre Channel FC-PH frames and Sequences of frames to implement the SCSI Fibre Channel Protocol (FCP).

1 Scope

This standard defines the Fibre Channel Protocol (FCP). The FCP is the Upper Layer Protocol ULP for applying the SCSI command set to the Fibre Channel.

2 Normative references

The following standards contain provisions which, through reference in the FCP, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards, and ANSI maintains registers for American National Standards.

1 **SCSI-2** - Small Computer System Interface - 2, draft, proposed ANSI Standard, X3T9.2/86-109, Revision 10H

2 **FC-PH** - Fibre Channel Physical and Signaling Interface, draft, proposed ANSI Standard, X3T9/91-062, Revision 3.0

3 **FC-FG** - Fibre Channel Fabric Requirements, draft requirements for proposed ANSI Standard, FC-FG/92-001, Revision 1.5

4 **SCSI CAM** - SCSI Common Access Method, draft, proposed ANSI Standard, X3T9.2/90-186, Revision 2.5

3 Definitions and conventions

Definitions, conventions, abbreviations, acronyms and symbols applicable to this standard are provided, unless they are identical to that described in FC-PH or SCSI-2, in which case they are included by reference.

3.1 Definitions

3.1.1 Base Address: a data buffer address meaningful to the Exchange Originator. In general, there can be one Base Address per Information Category but the FCP uses a Base Address only for the SCSI_DATA Information Category. The Relative Offset field in all frame headers of this Information Category indicates a displacement of the Payload with reference to the Base Address.

3.1.2 Fully Qualified Exchange ID (FQXID): an 80 bit token used to identify an I/O Process formed by the Source_ID || Destination_ID || Originator X_ID || Responder X_ID.

3.2 Editorial Conventions

In this standard, certain conditions, mechanisms, sequences, parameters, events, and similar terms are named and referred to by capitalizing the word in accordance with the conventions for proper names. Any lowercase uses of these words have the normal technical English meanings.

In all of the figures, tables, and text of this document, the most significant bit of a binary quantity is shown on the left side.

4 Structure and Concepts

The Fibre Channel (FC) is logically a point-to-point serial data channel. The architecture has been designed so that it can be implemented with high performance hardware that requires little real-time software management. The Fibre Channel Physical layer (FC-PH) performs all those functions required to transfer data from one node to another. The FC-PH can be treated as a very powerful delivery service with a three-layer hierarchy of information grouping and three defined classes of service.

The FC-4 protocol layers use the FC-PH signaling protocol to transfer the information required to execute the steps of an Upper Layer Protocol (ULP). The protocol is described in terms of the elements visible in the stream of Fibre Channel frames, Sequences, and Exchanges generated by a pair of nodes that support the ULP.

The detailed implementation that supports that stream of frames is not defined, although initiator and responder nodes are assumed to have a common interface for use by all ULPs that is similar in characteristics to the service interface defined in Annex R of the FC-PH[2]. An interface similar in nature to the SCSI Common Access Method[4] is expected to manage communication between the operating system of the host node and the FCP Upper Layer Protocol.

A SCSI I/O Process is started when host software provides a SCSI CAM Control Block to the software interface of the FCP. The CAM Control block contains

all the information necessary for execution of a SCSI command, including the address and characteristics of data to be transferred by the command. The FCP then performs a series of operations using the FC-PH services to control the proper sequencing of Exchanges and Sequences required to execute the SCSI FCP. The Sequences that are returned by the Responder complete the exchange of data and status required by the SCSI FCP. When the SCSI FCP operation is completed, returned information is used to prepare and return the CAM Control Block to the software that requested the operation. The returned status indicates whether or not the I/O Process was successful. The completion of the I/O Process indicates that the SCSI device performed the desired operations with the transferred data and that the information was successfully transferred to or from the SCSI Initiator's memory. A protocol is provided to present error information if the I/O Process is not successful.

An arbitrary number of I/O Processes may be active at one time, depending on the queueing capabilities of the particular SCSI node.

The FCP is designed to take full advantage of the multiplexing and shared bandwidth capabilities of FC-PH Class 2 operation. Fabric management and FC-3 frames are supported as specified by the referenced standards. Class 1 operation is considered a simplified subset of the normal FCP operation and is fully supported by the appropriate FC-PH Connection protocols. Intermix and out of order frame delivery are implementation options. Class 3 operation does not support the FCP, although Class 3 service may be provided on links carrying FCP information.

Both FC-PH and SCSI allow any node to be an originator and any node to be a responder. For SCSI I/O Processes between a host and a peripheral subsystem, the host normally takes on the SCSI initiator role and the peripheral subsystem normally takes on the target role. For host to host or device to device communications, either node can take on the initiator role. For Asynchronous Event Notification, the peripheral device takes on the SCSI initiator role to inform the host, in its target role, that an asynchronous event has occurred.

4.1 Link Management

The FC-PH interface explicitly allows drivers above the interface to perform those link control frame Sequences required for the support of FC-3 functionality. The frames or primitive Sequences used by the FCP include:

Basic Requests:

- No Operation
- Abort Sequence
- Remove Connection
- Echo

Basic Replies:

- Basic Accept
- Basic Reject

Extended Requests:

- Login
- Logout
- Abort Exchange
- Read Connection Status
- Read Exchange Status Block
- Read Sequence Status Block
- Establish Streaming
- Estimate Credit
- Advise Credit
- Read Time-out Value
- Read Link Status

Extended Replies:

- Accept
- Link Application Reject

The F_CTL functions performed in support of standard Exchange and Sequence functions are implemented directly by the firmware and hardware that will support any FC-2 N_Port. The link control frames required for normal FC-PH Sequence and Exchange management are also built into the N_Port firmware and hardware. Such frames used by the FCP include:

F_CTL Frames
- Exchange/Sequence Control Frames
  (bits 23-16)
  Fill Data Bytes
  (bits 1-0)

Link Control Frames
- R_Rdy Primitive
- ACK_1
- ACK_N
- N_Port Busy (P_BSY)
- Fabric Busy (F_BSY)
- N_Port Reject (P_RJT)
- Fabric Reject (F_RJT)

The FCP shall support FC-PH Class 2 service. All the Class 2 functions shall be available and all the Class 2 rules shall apply.

[Editor's Note: While Class 1 and Class 1 intermix services perform the SCSI FCP correctly, some implementations may have undesirable performance characteristics. Should we explore or describe those characteristics in this document?]

5 FCP

This section defines the frames, Sequences, Exchanges and protocols of the FCP.

5.1 FCP Overview

The FCP is based on a two-level ULP paradigm. The I/O Process of the SCSI protocol is mapped into a ULP Exchange. The individual phases of a SCSI I/O Process that are required to execute the logical function expressed by a software interface similar to the SCSI CAM are mapped into Sequences. Those SCSI Phases that perform link management on the SCSI Bus are not implemented by the ULP. Link control is instead performed by FC-PH and FC-3 protocols and frame structures. This is explained by the following table.

Table 1:
SCSI and FCP
Corresponding Functions

SCSI function:	FCP function:
I/O Process	Exchange

Table 1:
SCSI and FCP
Corresponding Functions

SCSI function:	FCP function:
Phase	Sequence
Command Transfer	Command Sequence
Reconnection Pointers	Xfer Rdy Sequence
Data Transfer	Data Xfer Sequence
Status Transfer	Response Sequence

Only one Sequence shall be in process for a given Exchange at any time.

An FCP Sequence is composed of one or more frames carrying the required information plus whatever acknowledgments are required using the FC-PH protocol to complete and control the Sequence.

A Sequence, especially a SCSI_DATA Sequence, may be split into multiple Sequences depending on the maximum Sequence size negotiated by the FCP nodes. For SCSI_DATA Sequence, the maximum size will normally be established by the Mode Sense/ Mode Select command's Disconnect Reconnect Page parameters. Each SCSI_DATA Sequence shall be preceded by a SCSI_XFER_RDY Sequence which indicates the actual size of that sequence.

Addressability to the fiber channel N_Port is defined by the Source and Destination N_Port values. Identification of an I/O Process on the Fibre Channel (corresponding to the ITLQ nexus definition of SCSI) is achieved by using the Fully Qualified Exchange Identifier (FQXID) formed by referencing the Source_ID, the Destination_ID, the Originator X_ID, and the Responder X_ID. The method for identifying an I/O Process within a node is not defined by this standard. Typical implementations use a software pointer to the control block for the I/O Process as the unique software label for the I/O Process.

Addressability of SCSI Devices and Logical Units internal to an addressed subsystem is obtained through an Entity parameter provided in the Command Sequence. Subsequent identification of the I/O Process (Exchange) is done by using the FQXID. Management of the protocol is performed by the management of the

completion of individual Sequences and by proper passing of initiative, Link_Control, and Link_Data frames.

The initiation and ending of Exchanges and Sequences are controlled and indicated by the proper values in the F_CTL, SOF, and EOF fields of the frame header.

Tables2 and3 define the sequences used by the FCP and the characteristics of those sequences.

Table 2: SCSI SEQUENCES SENT TO TARGETS

#	PHASE	DATA BLOCK		F/M/L	SI	SC	RO	M/O	COMMENTS
		CAT.	CONTENT						
I1	CMD	2	COMMAND	F	T	0	0	M	SCSI CMD
I2	DATA	1	DATA	M	T	0	disp	M	ONLY DATA

Table 3: SCSI SEQUENCES SENT TO INITIATORS

#	PHASE	DATA BLOCK		F/M/L	SI	SC	RO	M/O	COMMENTS
		CAT.	CONTENT						
T1	DATA	5	XFR RDY	M	T	0	0	M	WRITE XFR RDY
T2	DATA	5	XFR RDY	M	H	0	0	M	READ XFR RDY
T3	DATA	1	DATA	M	H	0	disp	M	ONLY DATA
T4	STATUS	3	STATUS	L	T	0	0	M	STATUS BLOCK

Key:

#	Sequence Name
CAT.	Information Category of data block
CONTENT	Contents (Payload) of data block
F/M/L	First/Middle/Last sequence of exchange
SI	Sequence Initiative--Held or Transferred
SC	Sequence Count (SEQ_CNT)--Start from 0 or Continued
M/O	Mandatory/Optional Sequence
XFR RDY	Transfer Ready
RO	Relative Offset, Starts with 0/Starts with displacement from Base Address

5.2 FCP Frame Format

A frame uses the standard FC-PH defined synchronization, initiation, CRC, and termination characters defined by the FC-PH standard. The generic format of a FCP frame is as follows:

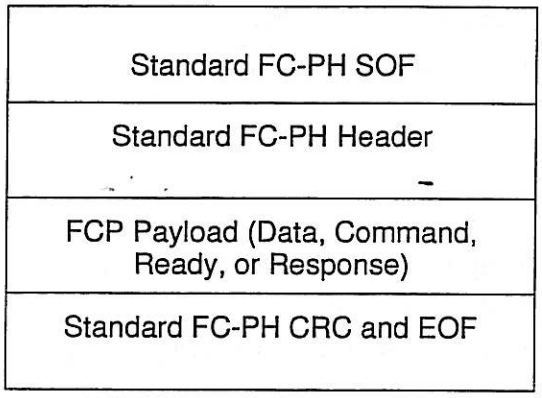

Figure 1
FCP Frame Format

The SOF and EOF delimiters should be used as specified by the FC-PH2. The high degree of parallelism characteristic of SCSI makes Class 2 service desirable.

SOF

FCP SOF delimiters operating on a Class 2 service use SOFi2 (Start_of_Frame Initiate Class 2) and SOFn2 (Start_of_Frame Normal Class2) frames. Other services use the appropriate SOF delimiters.

EOF

The FCP EOF delimiters shall use EOFt (End_of_Frame Terminate) delimiters for the last frame in a Sequence and EOFn (End_of_Frame Normal) delimiters for normal frames.

The FCP shall use the EOFni (End_of_Frame Invalid) delimiter to end a frame whose content was known to be invalid. FCP recipients of such a frame shall ignore the contents of the frame.

The FCP shall use the EOFa (End_of_Frame Abort) delimiter to end a frame which is known to have been damaged by the link. FCP recipients of such a frame shall ignore the contents of the frame.

Other frame ending delimiters may be called upon by services other than the FCP.

5.2.1 FCP Frame Header

The format of the standard FC-PH header as used by the FCP is as follows:

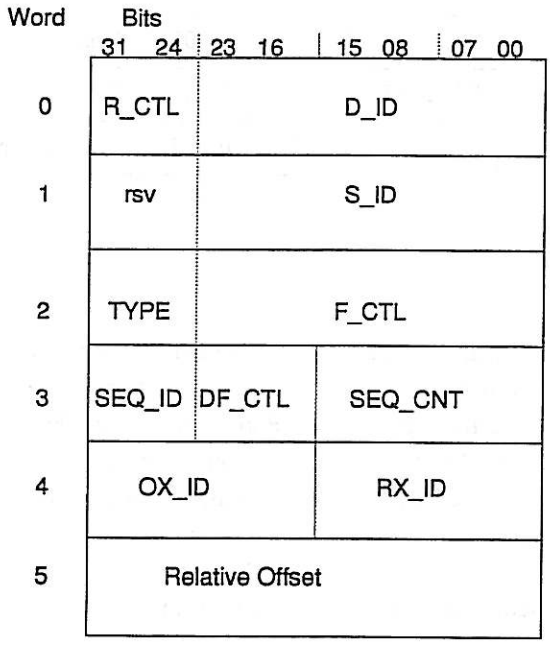

Figure 2
FCP Frame Header

All fields use the standard FC-PH definitions. The following explanations of some of the fields provide additional information about the field's usage in the FCP.

R_CTL

The R_CTL field identifies a particular frame type and Information Category. The Link_Data and Link_Control type frames are used by the FCP to perform standard FC-PH link control activities. All frames containing FCP information content shall be Device_Data type frames.

The Device_Data frame Information Categories are defined for the FCP frames and sequences in tables 2 and 3. The Information Categories used by the FCP are defined in Table 4.

Table 4: FCP Device_Data frame Information Categories

Category	Value	FC-PH Name	Description
SCSI_DATA	(0x01)	Solicited Data	Frames contain SCSI Data
SCSI_CMND	(0x02)	Unsolicited Control	Frames contain SCSI Command Block
SCSI_RSP	(0x03)	Solicited Control	Frames contain SCSI Response Block
SCSI_XFER_RDY	(0x05)	Transfer Ready	Frames contain SCSI Transfer Ready Block

D_ID

The D_ID identifies the destination of the frame. The D_ID may be the address of the N_Port or may address an alias for a group of N_Ports, any one of which may respond. The D_ID is one segment of the FQXID and is used unchanged in subsequent frames belonging to the same Exchange. If the D_ID is an alias, the Exchange uses the facilities made available to that alias, which may include hunt groups, dynamic path reconnection, or striping.

S_ID

The S_ID identifies the source of the frame. The S_ID may be the address of the Source N_Port or may be the address of an alias for a group of related N_Ports that constitute a source. The S_ID is one segment of the FQXID and is used unchanged in subsequent frames belonging to the same Exchange. A frame returned to an alias may be accepted by any of the N_Ports, depending on the facilities established for the alias.

TYPE

The TYPE field shall be 0x08, indicating this is a FCP frame. Because of SCSI's peer-to-peer structure, the 0X09 code, SCSI Target, will not be used and will become a reserved code.

F_CTL

The F_CTL (Frame Control) field manages the beginning and normal/abnormal termination of Sequences and Exchanges. The bits and definitions shall be as defined by the FC-PH, with the exception that the following Control Field bits shall have the required values indicated below.

End_Connection: Always 0. No Class 1 Connections are formed.

Connection_Resource: Always 0. No Class 1 Connections are formed.

New X_ID Assigned: Always 0. The FCP always retains the initial X_ID for the duration of an Exchange.

Invalidate X_ID: Always 0. The FCP always retains the initial X_ID for the duration of an Exchange.

[Editor's Note: The above restrictions are still under discussion. The best performance and simplest implementations will be generated using these assumptions, however:

a) Class 1 service would probably operate with no modifications to the ULP.

b) Systems tolerant of FQXID's modified by invalidating the X_ID would probably operate with no modifications to the ULP.]

SEQ_ID

The SEQ_ID as defined by the FC-PH uniquely identifies the frame as belonging to a particular Sequence.

DF_CTL

The DF_CTL is 0x00 to indicate that there are no

optional headers present.

[Editor's Note: The above restriction is still under discussion. The best performance and simplest implementation will be generated using that assumption.]

SEQ_CNT

The SEQ_CNT field indicates the frame number within the Sequence as defined by the FC-PH.

OX_ID

The OX_ID field is the Originator identification of the Exchange and is an element of the FQXID.

RX_ID

The RX_ID field is the Responder identification of the Exchange and is an element of the FQXID.

Relative_Offset

The Relative_Offset field indicates the relative displacement of the first byte of each frame's Payload with reference to the Base Address of the Information Category. See section 5.2.3. and 5.2.4

5.2.2 SCSI_CMND Sequence

Frames of the SCSI_CMND Sequence use the standard FC-PH frame header, described above. Because of the fixed short length of the SCSI_CMND Sequence's Payload, it is recommended that the Sequence consist of a single frame of information. The Sequence (frame) shall contain the following values and control fields in its Payload.

Table 5: FCP SCSI_CMND Payload

Field	Size
SCSI_ENT_ADDR	8 bytes
SCSI_CNTL	4 bytes
SCSI_DL	4 bytes
SCSI_CDB	16 bytes

SCSI_ENT_ADDR

The SCSI Entity Address is the mapping of the address of the desired device in the attached subsystem. The address space allows for any arbitrary drive/controller hierarchy, presently limited to 4 levels.

The hierarchy segments are separated into one 1-byte address component and three 2-byte address components. For each address component, a bit is available to indicate whether the component address is a LUN or Target address. The entity chosen selected by the address is the first LUN or target routine that is encountered while traversing the hierarchy. The addressed entity is always capable of responding to the SCSI command set.

Commands addressed to an addressed entity that does not exist are handled according to the SCSI rules for selection of Invalid Logical Units. The address may be identified as invalid at any level of the Entity Address.

Table 6: SCSI_ENT_ADDR Format

Field	Size
ENT_LUNTAR	1 byte
ENT_ADDR_0	1 byte
ENT_ADDR_1	2 bytes
ENT_ADDR_2	2 bytes
ENT_ADDR_3	2 bytes

ENT_ADDR_0 is always the first and highest layer of the hierarchy. The depth of the hierarchy of addressing, up to a maximum of four layers, is arbitrary and device dependent. The ENT_LUNTAR byte uses bits 0 (Most Significant Bit) through 3 to indicate in a bit significant manner which level of the hierarchy contains a Target Routine Identifier instead of a Logical Unit Number. The Target Routine bit must only be set for the lowest level of the hierarchical address applicable to the particular path.

SCSI_CNTL

The SCSI_CNTL field contains a number of control flags and control bits arranged in the following format.

Table 7: SCSI_CNTL Field Format

Bit Definition	Bit
Byte 0 (Most Significant)	
reserved	7-2
Read Data	1
Write Data	0
Byte 1	
reserved	7-3
Abort	2
Abort Tag	1
Reset	0
Byte 2	
reserved	7-0
Byte 3 (Least Significant)	
reserved	7
Simple_Q	6
Ordered_Q	5
Head_of_Q	4
reserved	3-0

Read Data - If this bit is set to one, expect the SCSI_DATA Sequence to be in the direction opposite to the direction of the Command Sequence (i.e., a read or inbound operation).

Write Data - If this bit is set to one, expect the SCSI_DATA Sequence to be in the same direction as the Command Sequence (i.e., a write or outbound operation). If both Read Data and Write Data are set to zero, there will be no SCSI_DATA Sequence.

Abort - If this bit is set to one, abort all outstanding I/O Processes for the addressed entity and all entities lower down in the hierarchy. A response Sequence will be returned for each aborted IO. The SCSI_CDB field will not be used.

Abort Tag - If this bit is set to one, abort the specified I/O Process. A response Sequence will be returned for an aborted IO. The Tag of the IO to abort (FQXID) will be in the SCSI_CDB field.

Reset - If this bit is set to one, reset the addressed entity and all entities lower down in the hierarchy.No response Sequences will be returned for any outstanding IOs. The SCSI_CDB field will not be used.

Simple_Q - If this bit is set to one, treat the FCP Exchange according to the rules for a Simple Queue Tag. If all three of the above queue management bits are zero, treat the FCP Exchange according to the rules for an untagged command.

Ordered_Q - If this bit is set to one, treat the FCP Exchange according to the rules for an Ordered Queue Tag.

Head_of_Q - If this bit is set to one, treat the FCP Exchange according to the rules for a Head of Queue Tag.

SCSI_DL

The SCSI_DL field contains a count of the total number of data bytes to be transferred in the SCSI data transfer Sequences for this SCSI I/O Process. The count is interpreted according to the context of the command. For devices with predictable transfer lengths, the count represents the number of bytes to be transferred. For devices with unpredictable transfer lengths, the count represents the maximum total transfer allowed for the I/O Process.

SCSI_CDB

The SCSI_CDB field contains the actual CDB to be interpreted by the SCSI Logical Unit addressed by the Entity Address. The maximum CDB length is constrained to be 16 bytes.

5.2.3 SCSI_XFER_RDY Sequence

The SCSI_XFER_RDY Sequence is composed of only one frame from the Responder followed by the appropriate acknowledgment frame from the Originator. The SCSI_XFER_RDY Sequence indicates that the Responder's buffer registers are allocated and prepared (full for a read, empty for a write) to perform all or a portion of the data transfer. The size of the expected data transfer is indicated by the Burst Length parameter.

A data transfer from or to the Responder may be divided by the Responder into more than one SCSI_DATA Sequence to meet its buffering requirements. Each SCSI_DATA Sequence must be preceded by the corresponding SCSI_XFER_RDY Sequence. The Originator is assumed to have available or be able to make available enough data buffer space to contain the entire requested data transfer with no more warning than the SCSI_XFER_RDY Sequence.

SCSI_XFER_RDY Sequences are required for both read and write (toward the Originator and from the Originator) SCSI_DATA Sequences. They serve the dual purpose of warning the Originator that a high speed Class 2 connection is about to resume while at the same time allowing the Responder to properly control its internal buffer structures.

Table 8: FCP SCSI_XFER_RDY Payload

Field	Size
SEQ_RO	4 bytes
BURST_LEN	4 bytes

SEQ_RO

The SEQ_RO field indicates the Relative Offset for frame 0 of the SCSI_DATA Sequence. This may be used by the Responder to deliver out of order Sequences on reads and to request out of order Sequences on writes.

BURST_LEN

The BURST_LEN field indicates the amount of buffer space prepared for the next SCSI_DATA Sequence. For most FCP environments and for typical SCSI data transfer lengths, this will be the entire data transfer.

5.2.4 SCSI_DATA Sequence

Frames in the SCSI_DATA Sequence use the standard FC-PH header. The data is identified with a particular SCSI I/O Process by the FQXID. The SEQ_ID and the Relative_Offset information is used to be sure that the frames are all received and placed in memory in the proper order. The entire data field of each frame is SCSI data and contains no optional headers.

The data to be transmitted during an I/O Process consists logically of one contiguous set of data whose first byte is at the Base Address. Data from that contiguous set of data may be requested in any order by the SCSI_XFER_RDY Sequence. The SEQ_RO paramter of the SCSI_XFER_RDY Sequence indicates the displacement from the Base Address for the first data of a particular SCSI_DATA Sequence. When the SCSI_DATA Sequence is transferred, frame 0 of the sequence will have that Relative Offset and subsequent frames will have correspondingly higher Relative Offsets. Frames and sequences may be delivered out of order. Depending on the implementation, data from the contiguous set of data may actually be scattered or gathered across any arbitrary physical memory boundaries independently of frame boundaries and Sequence boundaries. Such scatter/gather is not described by this standard.

5.2.5 SCSI_RSP Sequence

The SCSI_RSP Sequence uses the standard FC-PH frame header. The frame's data field contains the following values and control fields. The SCSI_RSP Sequence is typically composed of only one inbound frame followed by an outbound frame that acknowledges and terminates the Sequence. The SCSI FCP requires by convention that the total Response Information returned by the SCSI_RSP Sequence shall not exceed 1024 bytes to facilitate buffer management. The information contained in the SCSI_RSP Sequence is indicated in the following table.

Table 9: FCP SCSI_RSP Payload

Field	Size
SCSI_STATUS	4 bytes

Table 9: FCP SCSI_RSP Payload

Field	Size
SCSI_RESID	4 bytes
SCSI_SNS_LEN	4 bytes
SCSI_RSP_LEN	4 bytes
SCSI_SNS_INFO	n bytes
SCSI_RSP_INFO	m bytes

SCSI_STATUS

The SCSI_STATUS field is normally zero upon successful completion of a SCSI I/O Process. A zero value means no error and in addition indicates that no other information is present. A nonzero value means that there was an error, or that there is sense or response information included in the SCSI_RSP.

Bytes 0-1 are reserved and set to zero.

Byte 2 is nonzero if SCSI_RESID is nonzero, or if SCSI_SNS_LEN is nonzero, or if SCSI_RSP_LEN is nonzero.

Byte 3 (LSB) contains the status byte from the SCSI Logical Unit. The status byte codes are defined by the SCSI standard.

SCSI_RESID

The SCSI_RESID field contains a count of the number of residual data bytes which were not transferred in the SCSI data transfer Sequences for this SCSI I/O Process. A SCSI_RESID of zero is normal upon successful completion of a SCSI I/O Process. Devices having indeterminate data lengths may have a nonzero residual byte count for valid operations.

SCSI_SNS_LEN

The SCSI_SNS_LEN field specifies the number of valid bytes of SCSI_SNS_INFO. The number shall be an integral multiple of 4. A SCSI_SNS_LEN of zero specifies that no sense information is being provided. A SCSI_SNS_LEN of zero is normal but not required upon successful completion of a SCSI I/O Process.

SCSI_RSP_LEN

The SCSI_RSP_LEN field specifies the number of valid bytes of SCSI_RSP_INFO. The number shall be an integral multiple of 4. A SCSI_RSP_LEN of zero specifies that no response information is being provided. a SCSI_RSP_LEN of zero is normal but not required upon successful completion of a SCSI I/O Process.

SCSI_SNS_INFO

The SCSI_SNS_INFO contains the information specified by the SCSI-2 Standard for presentation by the Request Sense command. The SCSI-2 Standard requires that a minimum of 18 bytes be presented if any bytes are presented. The proper SCSI_SNS_INFO shall be presented when the SCSI Status byte of Check Condition or Busy is presented. The SCSI_SNS_INFO is optionally presented for any other SCSI Status byte values.

SCSI_RSP_INFO

The SCSI_RSP_INFO field contains information describing the completion of a SCSI I/O Process. All bytes of the field are reserved, pending definition.

6 FCP Protocol

6.1 FCP Typical Protocols

6.1.1 SCSI Write Example

All frames have a frame level acknowledgment understood to be returned automatically as part of the link control. The frame level acknowledgment uses the R_RDY primitive.

See figures 3 and 4 on the following pages.

6.1.2 SCSI Read Example

All frames have a frame level acknowledgment understood to be returned automatically as part of the link control. The frame level acknowledgment uses the R_RDY primitive.

See figures 5 and 6 on the following pages.

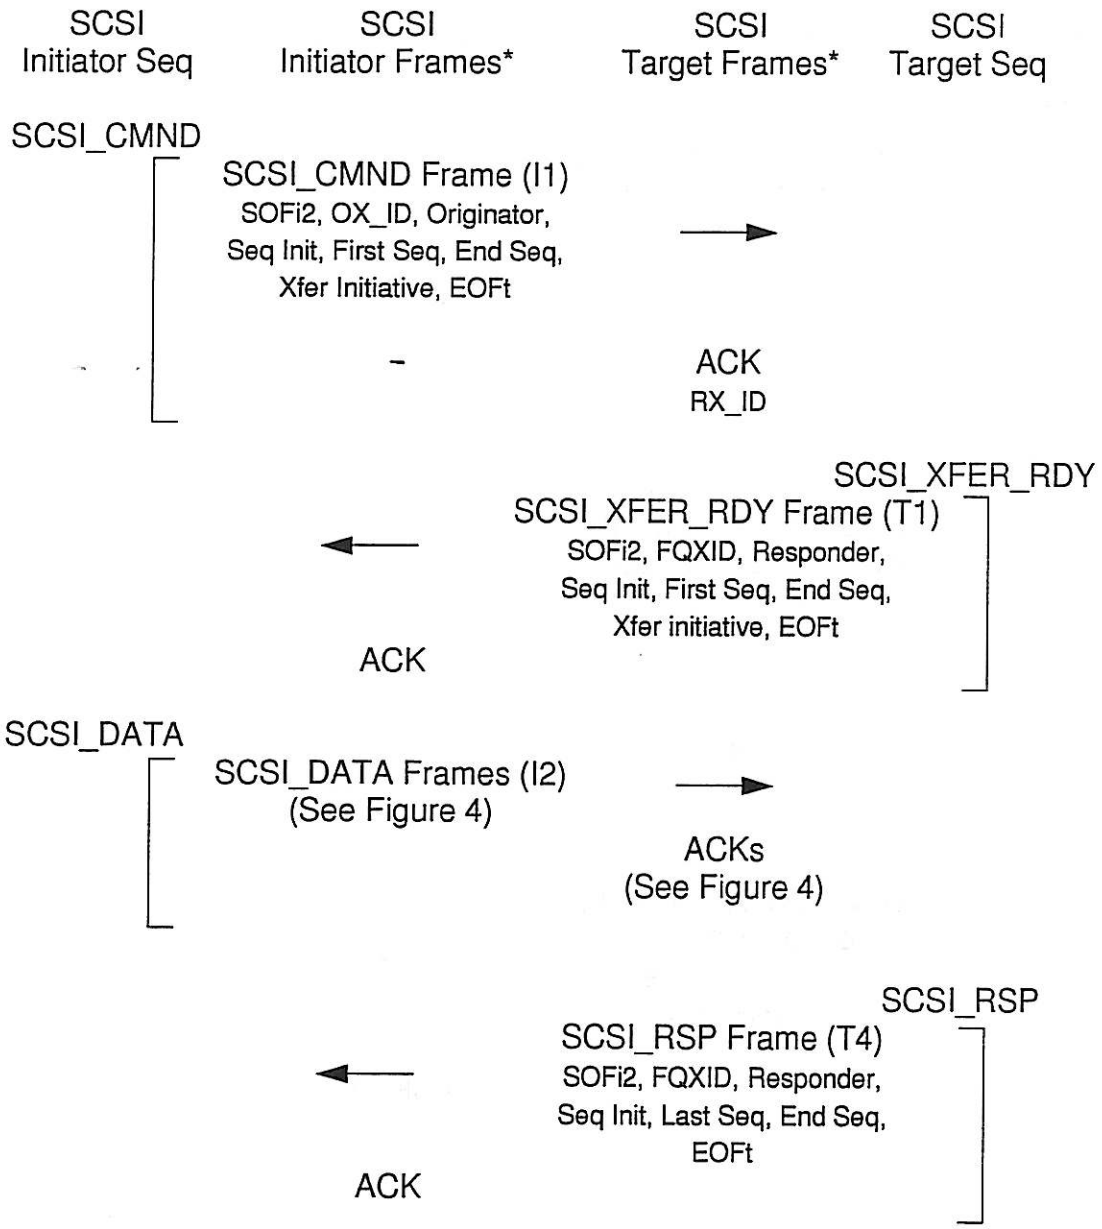

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 3
Example of SCSI Write I/O Process

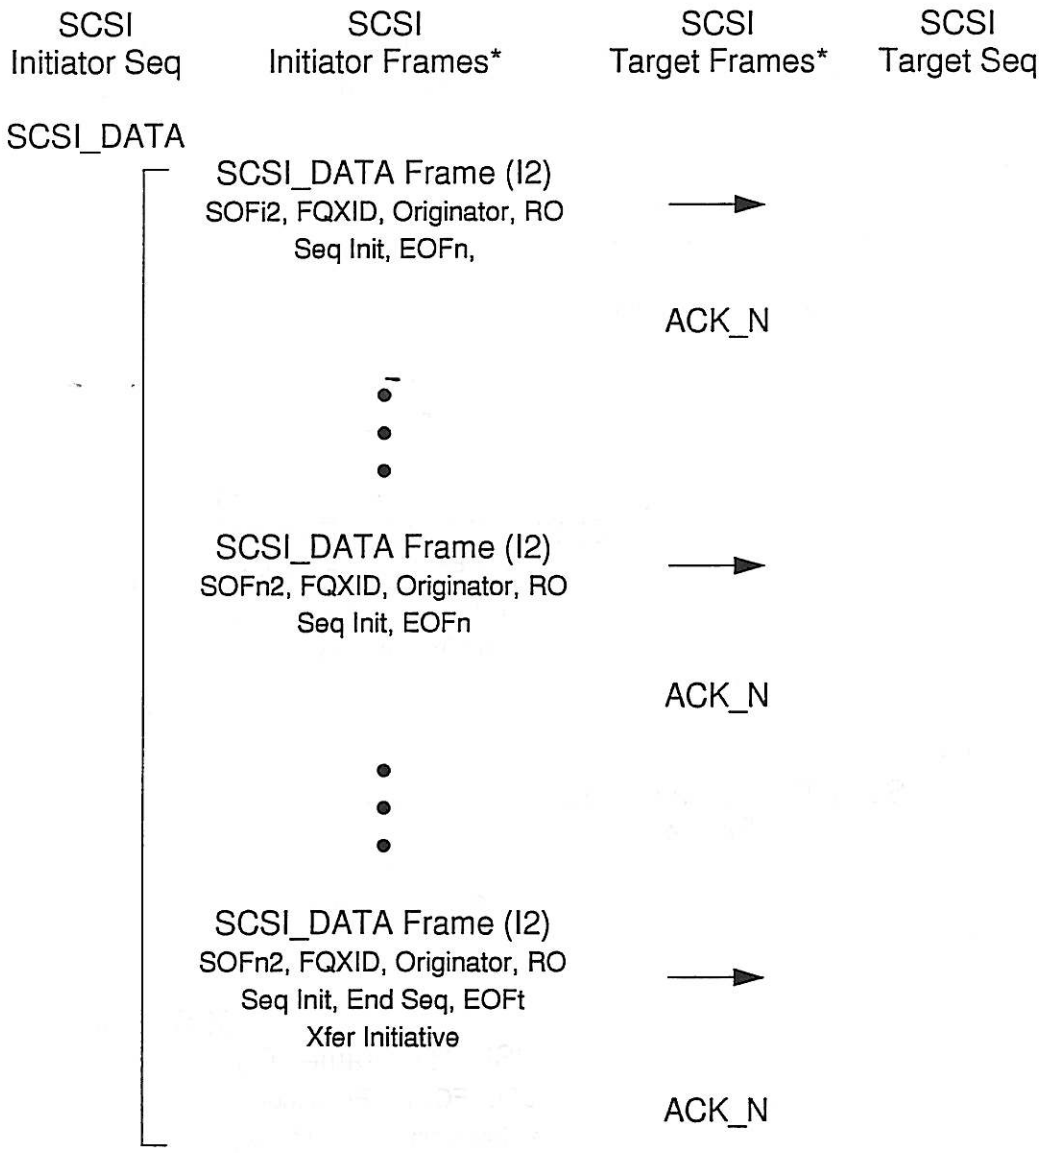

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 4
Example of SCSI_DATA Write Sequence

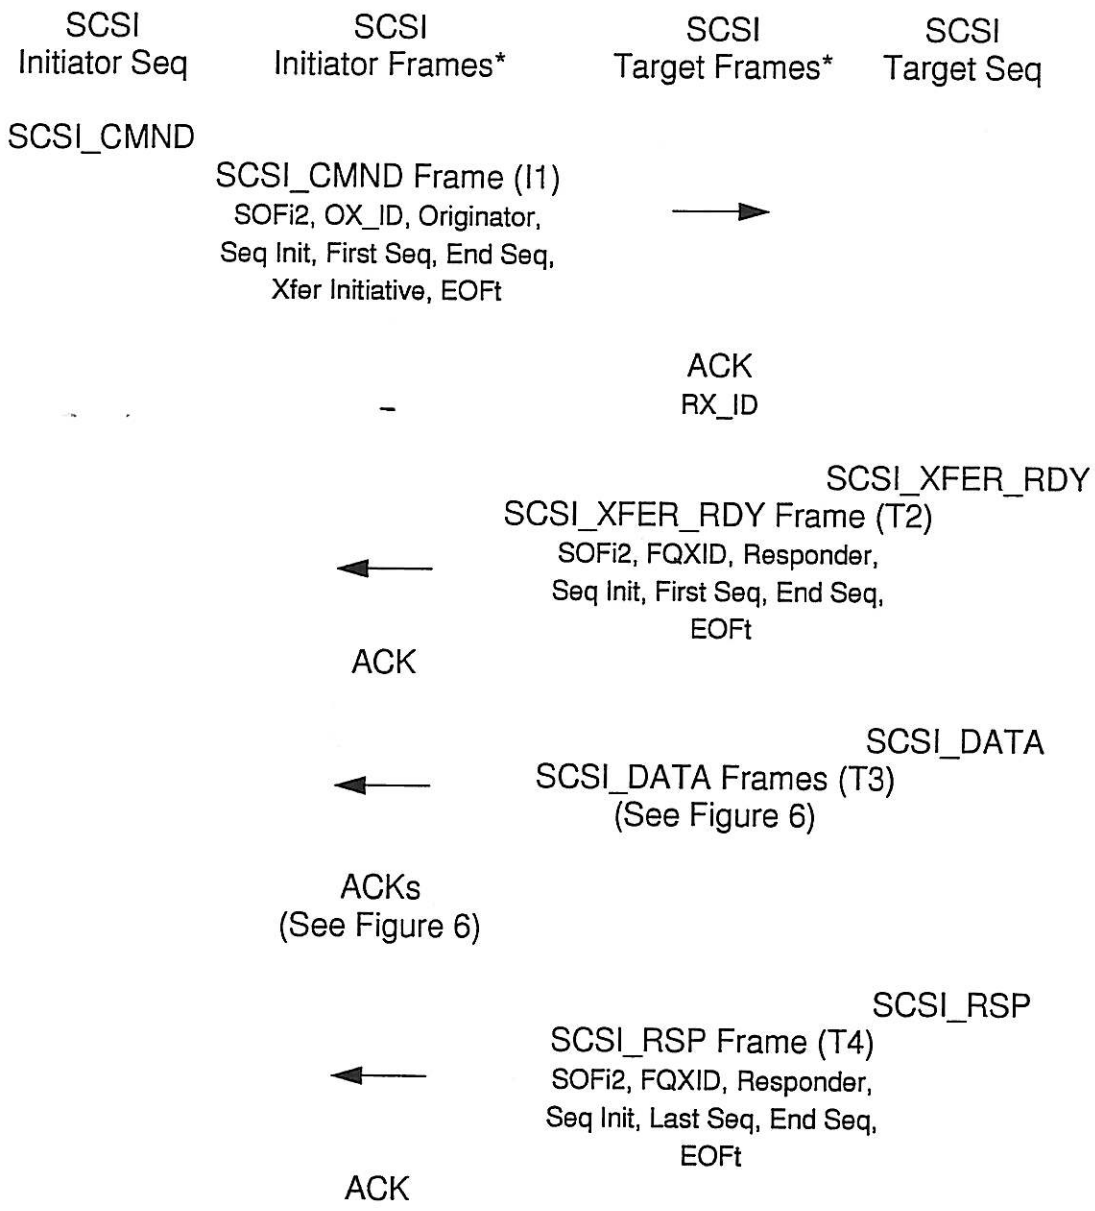

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 5
Example of SCSI Read I/O Process

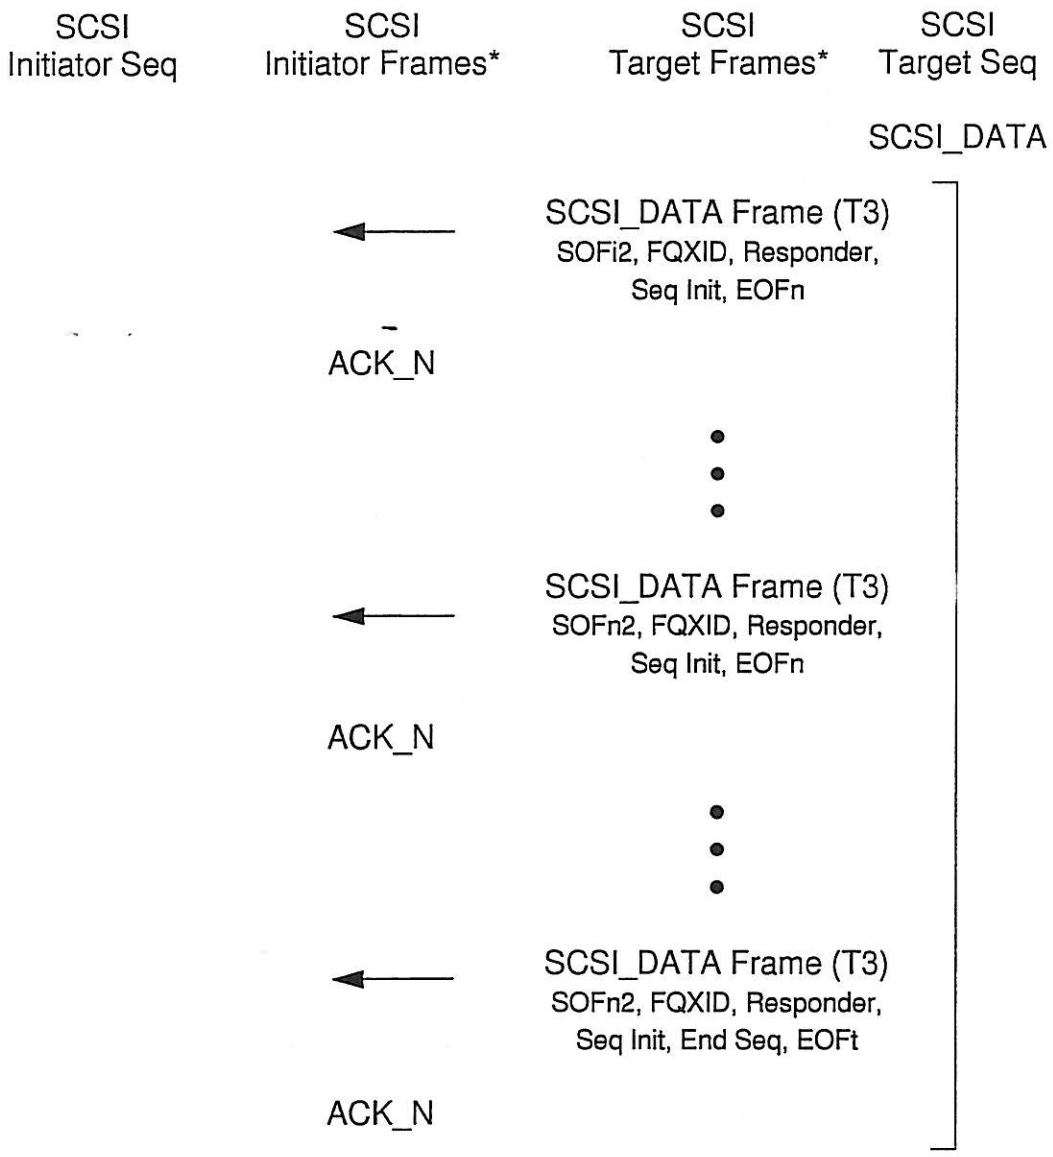

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 6
Example of SCSI_DATA Read Sequence