Sun Microsystems Proposal

FC-P/91-0__ R1.1
X3T9.2/91-122 R1.1
X3T9.3/91-____ R1.1

FIBRE CHANNEL

SCSI PROTOCOL (FC-4S)

REV 1.1

working draft proposed
American National Standard
for Information Systems
February 13, 1992

Secretariat:

Computer & Business Equipment Manufacturers Association

ABSTRACT: This standard 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 an internal 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@wichitaks.ncr.com

John Lohmeyer (X3T9.2 Chairman)
NCR Corporation
3718 N. Rock Road
Wichita, KS 67226
(316) 636-8703
E-Mail:

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-2 frames and Sequences of frames to implement the SCSI FC-4 (FC-4S).

Annex A describes the functions from the SCSI Standard, FC-F Standard and FC-3 proposals which are required to fully implement and support FC-4S. It is expected that all these services will also be required for other FC-4 protocols.

Annex B indicates which functions and options of FC-2 and SCSI-2 are not required for support of FC-4S.

1 Scope

This standard defines the Fibre Channel FC-4 implementation for the SCSI command set.

2 Normative references

The following standards contain provisions which, through reference in FC-4S, 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.

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

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

FC-PH - Fibre Channel Physical and Signaling Interface, draft, proposed ANSI Standard, X3T9.3/92-007

FC-F - Fibre Channel Fabric, draft requirements for proposed ANSI Standard

IPI FC-4 - Fibre Channel Implementation for IPI, not yet published

HIPPI FC-4 - Fibre Channel Implementation for HIPPI, not yet published

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 Data Category but FC-4S uses a Base Address only for the SCSI_DATA Data Category. The Relative Offset field in all frame headers of this Data 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 IO Process formed by the Source_ID || Destination_ID || Originator X_ID || Responder X_ID.

3.1.3 Request ID: a 32 bit token used to identify an IO Process. This token is unique within a host, but not unique across the Fabric.

3.2 Editorial Conventions

In this standard, a number of conditions, mechanisms, sequences, parameters, events, English text, states, or similar terms are printed with the first letter of each word in uppercase and the rest lowercase (e.g. Exchange, Class). Any lowercase uses of these words have the normal technical English meanings.

Numbered items in this standard do not represent any priority. If prioritized, it will be specifically so indicated.

The American convention of numbering is used, i.e., the thousands and higher multiples are separated by a comma and a period is used as the decimal point. This is equivalent to the ISO convention of a space and a comma.

American	ISO
0.6	0,6
1,000	1 000
1,323,462.9	1 323 462,9

In case of any conflict between figure, table, and text, the text takes precedence.

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. Exception to this conventions is indicted at the appropriate section.

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 logically operate devices of a particular type. The protocol is described in terms of the elements visible in the stream of frames, Sequences, and Exchanges generated by a pair of nodes that are supporting FC-4S.

The detailed implementation that supports that stream of frames is not defined, although initiator nodes are assumed to have an interface similar in nature to the SCSI Common Access Method (CAM) to manage communication with the FC-4/FC-2 software and hardware. A SCSI IO Process is started when software provides a CAM Control Block to the CAM interface. 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. An arbitrary number of IO Processes may be active at one time, depending on the queueing capabilities of the particular SCSI ITL nexus. When the SCSI IO Process is completed, the CAM Control Block with associated completion status is returned to the SCSI device driver that requested the operation. The returned status indicates whether or not the IO Process was successful. The completion of the IO 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 IO Process is not successful.

FC-4S is designed to take full advantage of the multiplexing and shared bandwidth capabilities of FC-2 Class 2 operation. Fabric management and FC-3 frames are supported as specified by the referenced standards, except as noted in Annex B. Class 1 operation is considered a simplified subset of the normal FC-4S operation and is fully supported by the appropriate FC-2 Connection protocols. Intermix and out of order frame delivery are implementation options. Class 3 operation does not support FC-4S.

4.1 Link Management

The FC-2 interface explicitly allows drivers above the interface to perform those link control frame Sequences required for the support of FC-3 functionality. Such frame or primitive Sequences used by FC-4S include:

Link_Data Requests
- Login
- Logout
- Abort Exchange
- Abort Sequence
- Read Connection Status
- Read Exchange Status Block
- Read Sequence Status Block
- NOP
- Estimate Credit
- Advise Credit
- Remove Connection

Link_Data 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-2 Sequence and Exchange management are also built into the N_Port firmware and hardware. Such frames used by FC-4S include:

F_CTL Frames
- Exchange/Sequence Control Frames (bits 23-16)

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)

FC-4S shall support FC-PH Class 2 operation. All the Class 2 functions shall be available and all the Class 2 rules shall apply.

5 FC-4S

This section defines the frames, Sequences, Exchanges and protocols of FC-4S.

5.1 FC-4S Overview

FC-4S is based on a two-level FC-4 paradigm. The IO Process of the SCSI protocol is mapped into a FC-4 Exchange. The individual phases of a SCSI IO 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 FC-4 protocol. Link control is instead performed by FC-2 and FC-3 protocols and frame structures. This is explained by the following chart:

SCSI function:	FC-4S function:
IO Process Phase Async Event	Exchange Sequence Exchange
Command Transfer Reconnection Pointers Data Transfer Status Transfer	Command Sequence Xfer Ready Sequence Data Xfer Sequence Response Sequence

Figure 1
FC-4S Sequence Definitions

Asynchronous events are mapped into Exchanges with standard FC-4S Sequences, but the Originator and Responder roles are reversed.

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

A Sequence, especially a data transfer Sequence (SCSI_DATA Sequence), shall be composed of one or more frames carrying the required information plus whatever acknowledgments are required using the FC-2 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 specified by the appropriate SCSI_XFER_RDY Sequence. Each SCSI_DATA Sequence shall be preceded by a SCSI_XFER_RDY Sequence.

Addressability to the fiber channel N_Port is defined by the Source and Destination N_Port values. Identification of an IO process (corresponding to the ITLQ nexus definition of SCSI) is achieved by using the Fully Qualified Exchange Identifier (FQXID) formed by the Source_ID||Destination_ID|| Originator X_ID||Responder X_ID. A Request_ID is used to relate a SCSI IO Process to the FQXID, its value is not specified by the standard, although there must be a one-to-one relationship.

Informative Note: As an example, the SCSI CAM typically associates the nexus of an IO Process with a pointer to the CAM Control Block.

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 IO 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.

5.2 FC-4S Frame Format

A frame uses the standard FC-2 defined synchronization, initiation, CRC, and termination characters defined by the FC-PH standard. The generic format of a FC-4S frame is as follows:

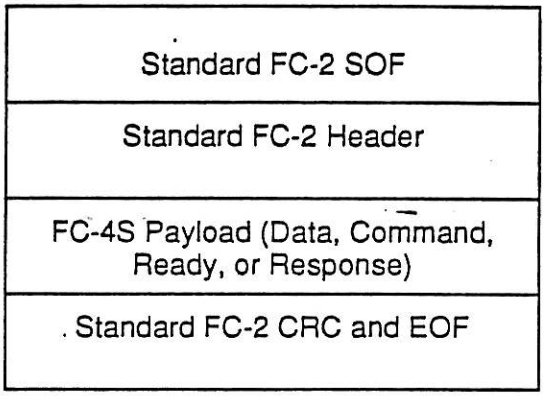

Figure 2
FC-4S Frame Format

SOF

FC-4S SOF delimiters shall use SOFi2 (Start_of_Frame Initiate Class 2) and SOFn2(Start_of_Frame Normal Class2) frames only. Services other than FC-4S may use any appropriate SOF delimiters.

EOF

FC-4S 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.

FC-4S shall use the EOFni (End_of_Frame Invalid) delimiter to end a frame whose content was known to be invalid. FC-4S recipients of such a frame shall ignore the contents of the frame.

FC-4S shall use the EOFa (End_of_Frame Abort) delimiter to end a frame which is known to have been damaged by the link. FC-4S recipients of such a frame shall ignore the contents of the frame.

Other frame ending delimiters may be called upon by services other than FC-4S.

5.2.1 FC-4S Frame Header

The format of the standard FC-2 header as used by FC-4S is as follows:

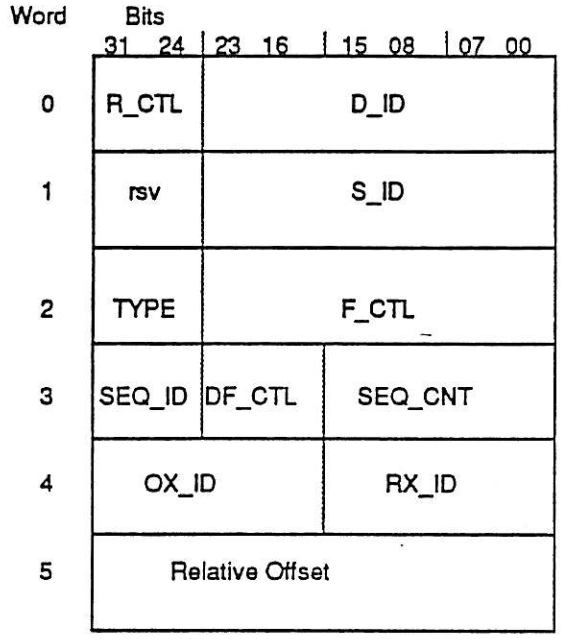

Figure 3
FC-4S Frame Header

R_CTL

The R_CTL field identifies a particular frame type and Data Category. The Link_Data and Link_Control type frames are used by FC-4S to perform standard FC-2 link control activities. All frames containing FC-4S information content shall be Device_Data type frames.

The following Device_Data frame Data Categories are defined for FC-4S.

SCSI_XFER_RDY Category (0x00) - Indicates frames contain SCSI Transfer Ready Block

SCSI_DATA Category (0x01) - Indicates frames contain SCSI Data

SCSI_CMND Category (0x02)- indicates frames contain standard SCSI Command Block

SCSI_RSP Category (0x03) - Indicates frames contain SCSI Response 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 FC-4S 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-2, with the exception that the following Control Field bits shall be a fixed state to properly implement FC-4S.

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. FC-4S always retains the initial X_ID for the duration of an Exchange.

Invalidate X_ID: Always 0. FC-4S always retains the initial X_ID for the duration of an Exchange.

SEQ_ID

The SEQ_ID as defined by the FC-PH uniquely identifies the frame as belonging to a particular Sequence. This is used to detect certain types of errors that may cause more than one Sequence to be active at a time in an Exchange.

DF_CTL

The DF_CTL is 0x00 to indicate that there are no optional headers present.

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 Data Category.

5.2.2 SCSI_CMND Sequence

Frames of the SCSI_CMND Sequence use the standard FC-2 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.

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

Figure 4
FC-4S
SCSI_CMND Sequence Payload

SCSI_ENT_ADDR

This is the mapping of the address of the desired device in the attached subsystem. The address space allows for any arbitrary drive/controller hierarchy.

The hierarchy segments are separated into one 1-byte address component and several 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 selected entity 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, no matter which levels of the hierarchy force the address to not match a valid SCSI_ENT_ADDR.

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

Figure 5
SCSI_ENT_ADDR Format

ENT_ADDR_3 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. If a bit in the ENT_LUNTAR field is set, it indicates that at the corresponding layer of the hierarchy, a Target Routine, rather than a LUN, will be addressed. The Target Routine bit must only be set for the lowest level of the hierarchical address applicable to the particular path.

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.

SCSI_CNTL

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

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

Figure 6
SCSI_CNTL Field Format

Abort - If this bit is set to one, abort all outstanding IO 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 IO 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.

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.

Head_of_Q - If this bit is set to one, treat the FC-4S Exchange according to the rules for a Head of Queue Tag.

Ordered_Q - If this bit is set to one, treat the FC-4S Exchange according to the rules for an Ordered Queue Tag.

Simple_Q - If this bit is set to one, treat the FC-4S Exchange according to the rules for a Simple Queue Tag. If all three of the above queue management bits are zero, treat the FC-4S Exchange according to the rules for an untagged command.

SCSI_DL

The SCSI_DL field contains a count of the total number of data bytes expected to be transferred in the SCSI data transfer Sequences for this SCSI IO Process.

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 SCSI Burst Size 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.

Field	Size
SCSI_BURST_LEN	4 bytes

Figure 7
FC-4S
SCSI_XFER_RDY Payload

SCSI_BURST_LEN

The SCSI_BURST_LEN field indicates the amount of buffer space prepared for the next SCSI_DATA Sequence. For most FC-4S environments and for typical SCSI data transfer lengths, this will be the entire data transfer.

Relative_Offset

The Relative_Offset field in the frame header of the SCSI_XFER_RDY Sequence defines the relative displacement of the first byte of the Payload of the following SCSI_DATA frames from the Base Address of the IO Process. This may be used by the Responder to deliver out of order Sequences on writes, and to request out of order Sequences on reads.

5.2.4 SCSI_DATA Sequence

Frames in the SCSI_DATA Sequence use the standard FC-2 header. The data is identified with a particular SCSI IO 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.

5.2.5 SCSI_RSP Sequence

The SCSI_RSP Sequence uses the standard FC-2 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. FC-4S requires by convention that the total Response Information returned by the SCSI_RSP Sequence shall not exceed 1048 bytes to facilitate buffer management. The information contained in the SCSI_RSP Sequence is indicated in the following figure.

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

Figure 8
FC-4S
SCSI_RSP Sequence Payload

SCSI_STATUS

The SCSI_STATUS field is normally zero upon successful completion of a SCSI IO Process. A zero value means no error and no other information is present. A nonzero value means there was an error, or there is sense or response information.

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.

Fibre Channel errors, such as P_BSY and P_RJT, are reported in the frame header for the SCSI_RSP Sequence. In this case the Sequence Payload is not valid.

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 IO Process. A SCSI_RESID of zero is normal upon successful completion of a SCSI IO Process..

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 IO 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 IO 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 IO Process. All bytes of the field are reserved, pending definition.

6 FC-4S Protocol

6.1 FC-4S Data Burst Management

End-to-end flow control management is used by FC-4S.

6.2 FC-4S Typical Protocols

6.2.1 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 figures10 and 11 on the following pages.

6.2.2 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 12 and 13 on the following pages.

6.2.3 SCSI Async Example

The SCSI Asynchronous Exchange is a Device Originated Exchange used for Asynchronous Event Notification.

See figure 14 on the following pages.

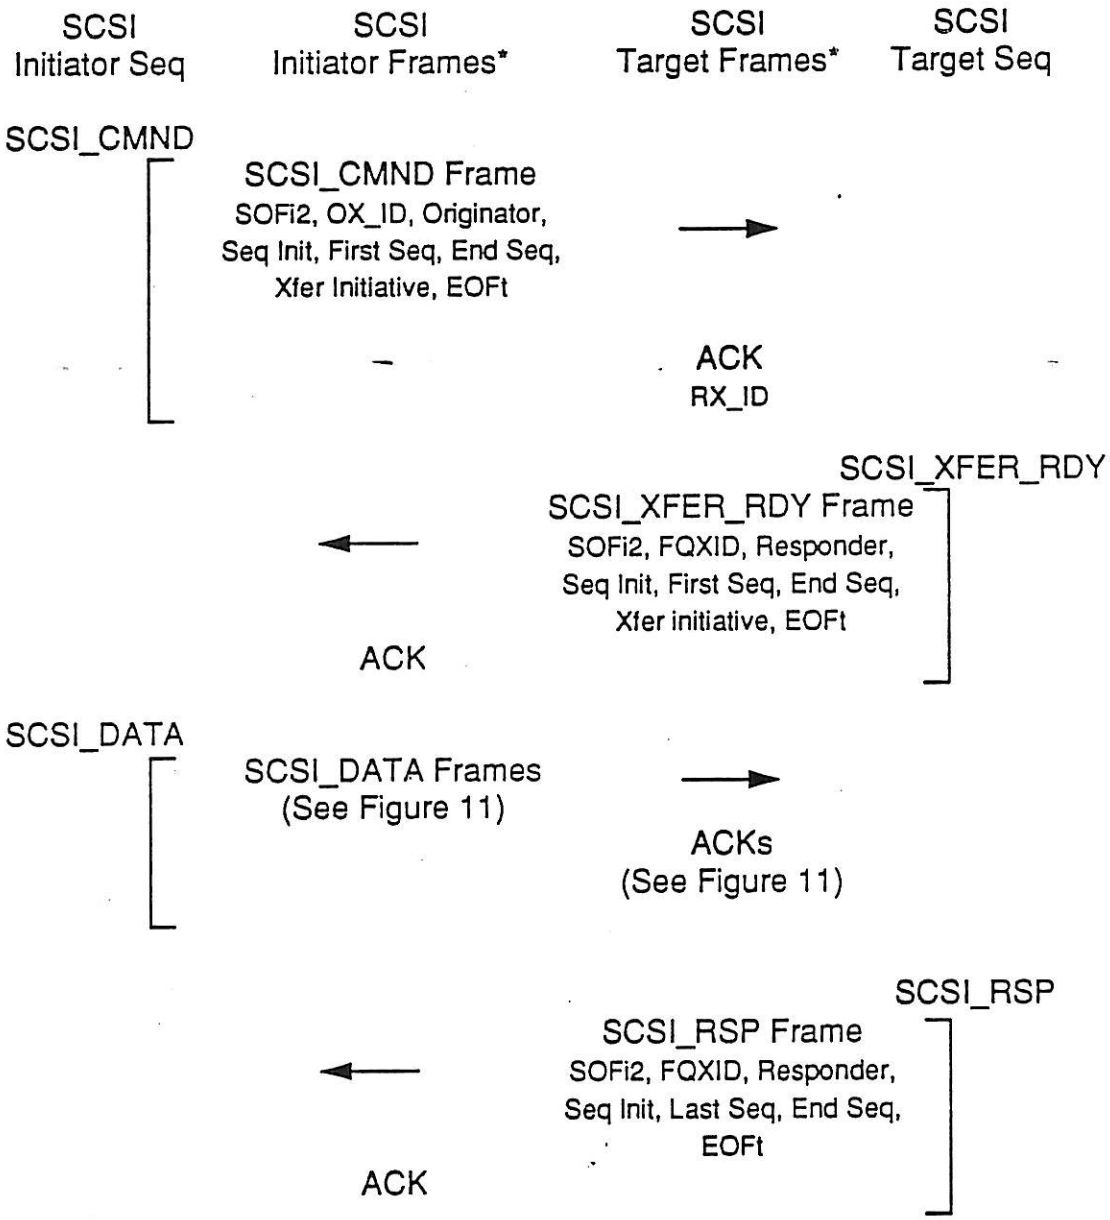

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 10
Example of SCSI Write IO Process

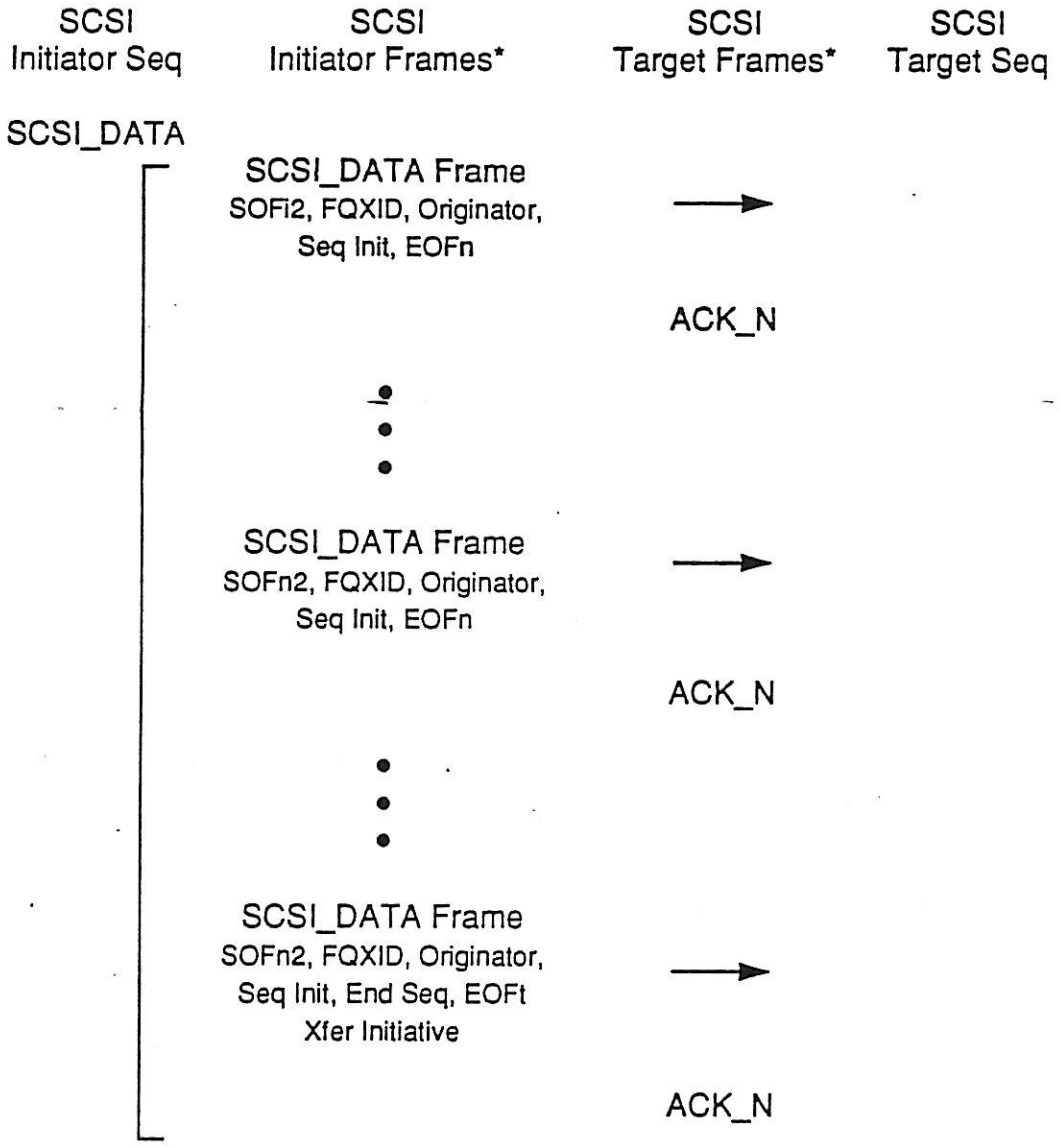

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 11
Example of SCSI_DATA Write Sequence

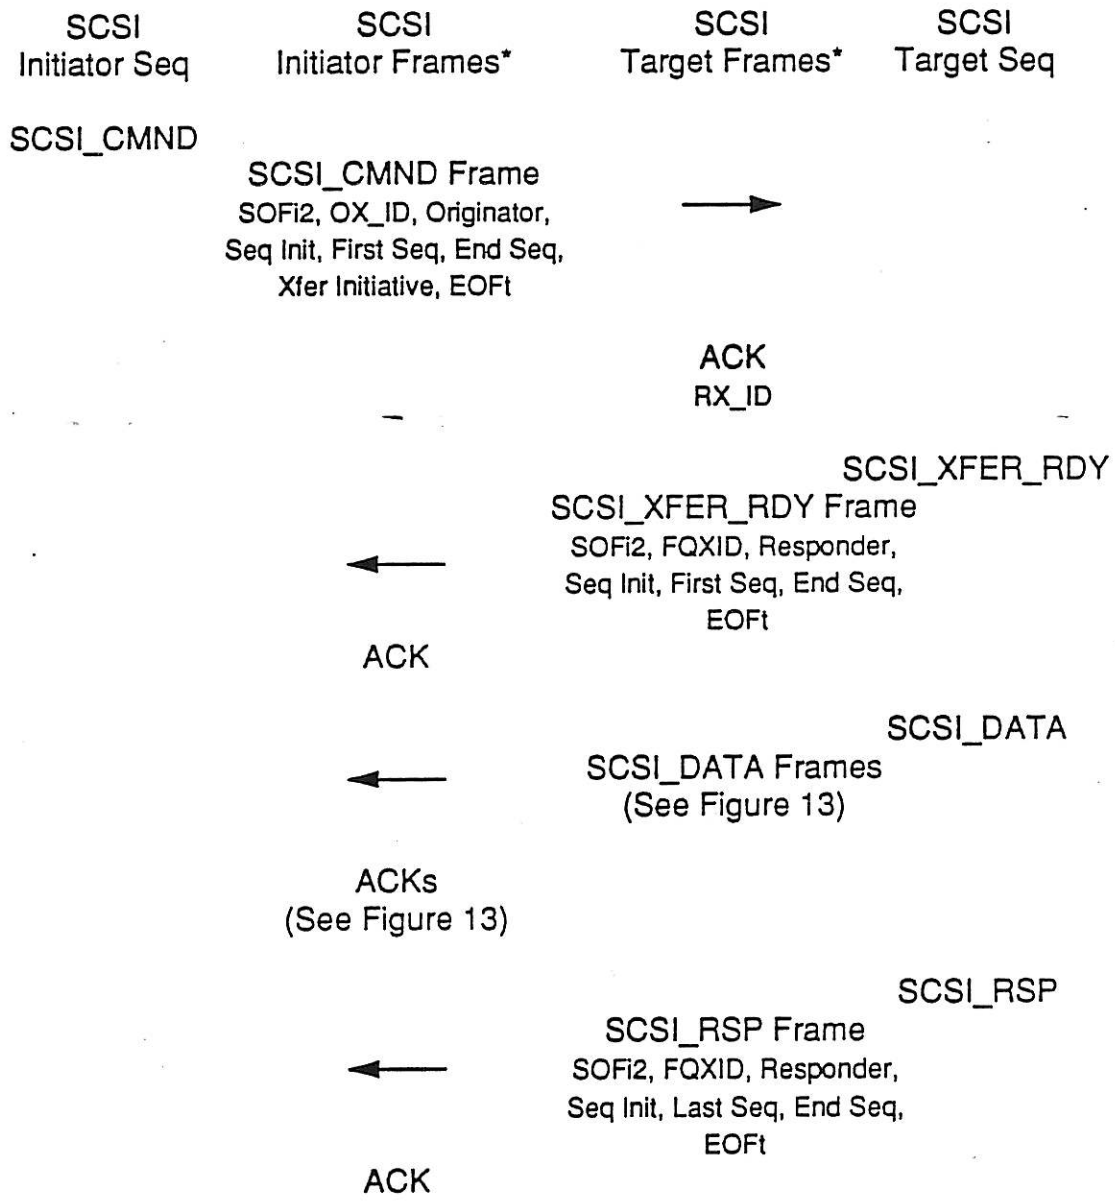

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 12
Example of SCSI Read IO Process

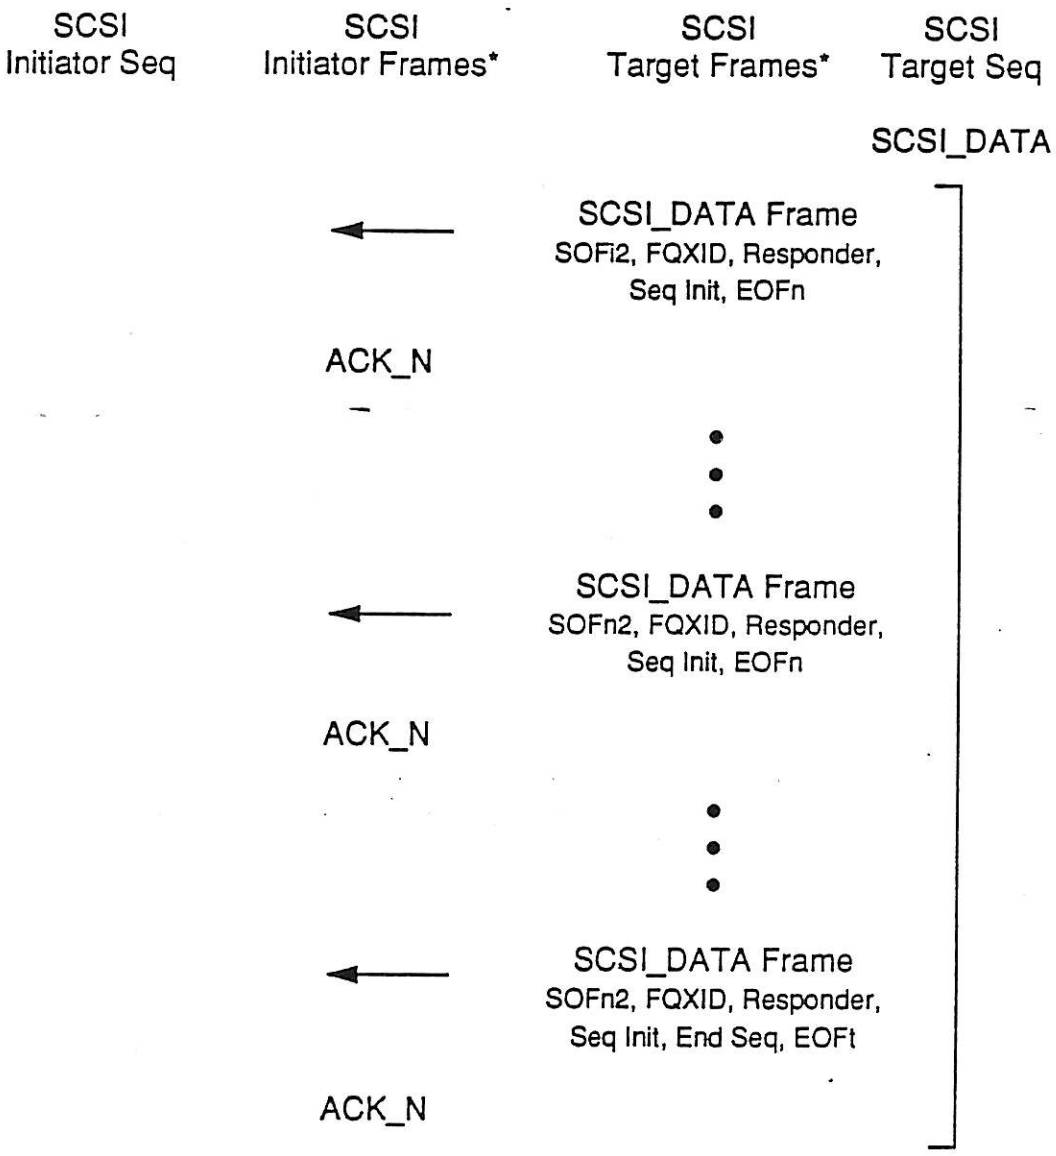

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 13
Example of SCSI_DATA Read Sequence

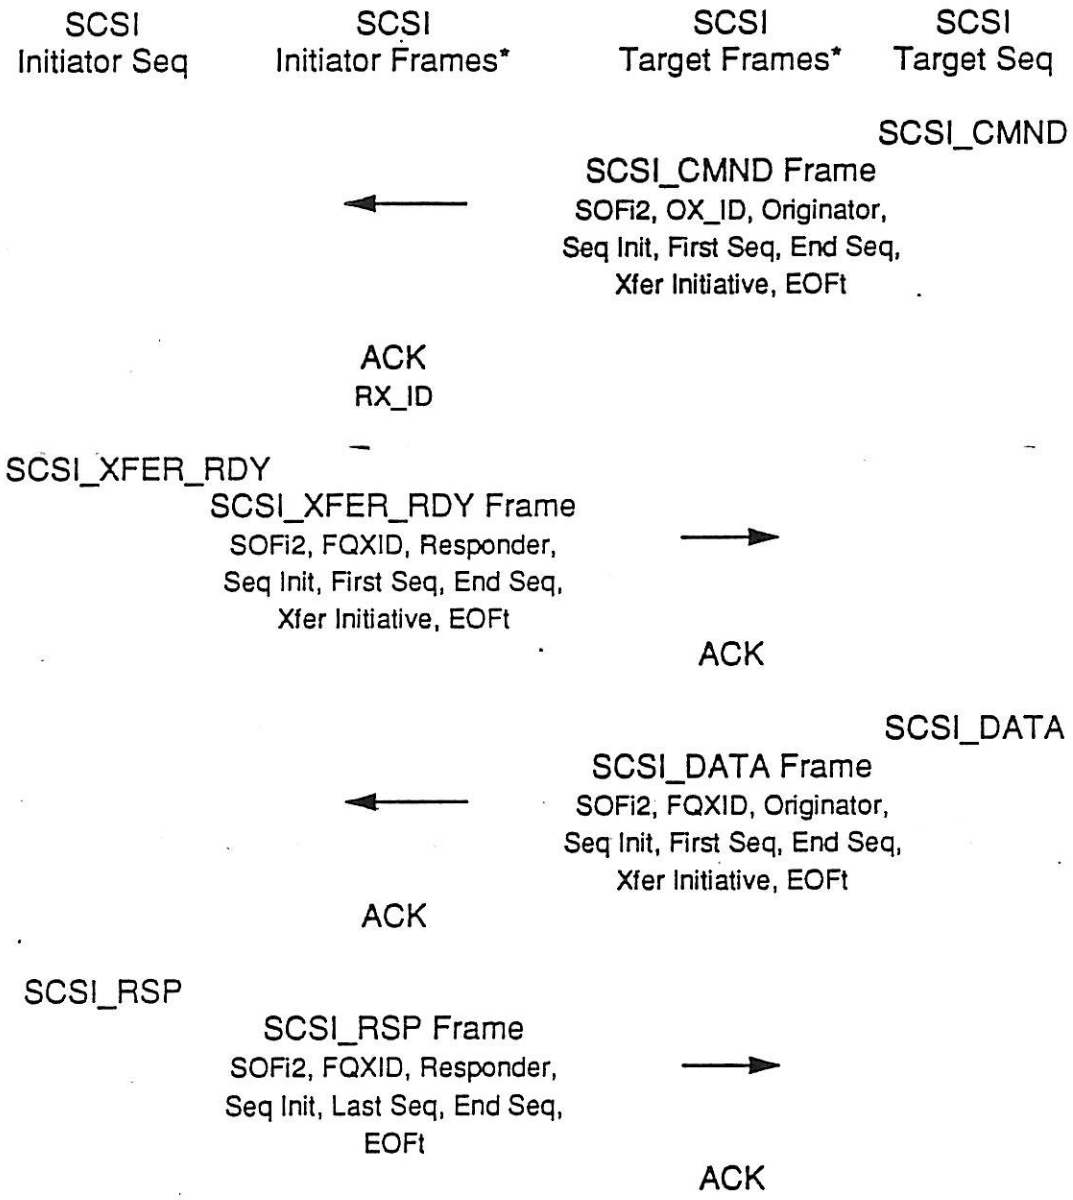

* Frame level acknowledgment using R_RDY is assumed as automatic part of frame transmission.

Figure 14
Example of SCSI Async Exchange

Annex A: Requirements

A.1 FC-2 Functionality Requirements

A.2 FC-3 Functionality Requirements

A.3 FC-F Functionality Requirements

A.4 SCSI Functionality Requirements

A.4.1 Configuration

A new Mode Sense page, "Configuration Definition" is required.

A.4.2 Entity Address

The Entity Address concept must be introduced to the proper level of SCSI. The LUN at the well known address 0xFF.FF|00.00|00.00|00.00 must be of a type FC-4S and must accept the Test Unit Ready, Request Sense, Inquiry, and Mode Sense commands.

Annex B: Unused Functionality

B.1 FC-2 Unused Functionality

B.1.1 TYPE Codes

The TYPE Code for SCSI is presently defined for both the initiator and target. The Code will be combined to a single code of 0x08, since the SCSI is a completely peer-to-peer interface whose Sequences are completely defined within an Exchange by the Data Category.

B.1.2 Optional Headers

No optional headers will be used by FC-4S.

B.1.3 Control Frame Data Category

Another Device_Data frame Data Category may be required if there are any SCSI functions not properly managed by already defined FC-2/FC-3 Link_Data and Link_Control frames. No such function is known at this time.

B.2 FC-3 Unused Functionality

B.3 FC-F Unused Functionality

B.4 SCSI Unused Functionality

B.4.1 CDB Length

The maximum length of a CDB in SCSI is not defined. FC-4S constrains the maximum length of a CDB to be 16 bytes. The standard CDB lengths are presently defined as 6, 10, and 12 bytes. Vendor unique CDB's may be defined, but 16 bytes provides room for expanding the present 10 byte CDB to 64 bit addressing and a 32 bit count field, adequate for most command sets.

Annex C: Concerns, Questions

C.1 Burst Throttling

Most architectures provide some kind of burst throttling. For best flexibility, a maximum burst size can be negotiated for each Sequence that is to be transferred. When a data transfer Sequence is broken up into multiple Sequences of the maximum burst size, it is possible to create more efficient use of second level buffer structures behind the primary FC-2 hardware buffers.

The maximum burst size is defined in the SCSI_XFER_RDY Sequence for the SCSI_DATA Sequence which is to follow the SCSI_XFER_RDY Sequence. The maximum burst size can be defined for either inbound or outbound transfer, although outbound transfer is the most difficult for the Responder to manage. It is assumed that Originator data buffers have already been allocated. It is assumed that the primary FC-2 hardware buffer throttling is managed by the credit definition.