

4 FCP/Class 2 Public Loop Operation for Streaming Devices (FCW/98w125r0,2/10/98)

This document describes required and optional behavior of SCSI streaming devices in a Public Loop configuration using FCP and class 2 delivery service. This document specifies operation from a FCP perspective. Refer to the FC-PH document(s) for rules and guidelines pertaining to class 2 operation.

Class 2 delivery service provides end-to-end flow control and sequence/frame level error detection and recovery. The environment defined by this document does not ensure in-order delivery of frames.

4.1 N_Port Class 2 Service Parameters - Fabric Login

Table 3 lists Class 2 Service Parameters with usage defined by this document. The following legend is used for entries in this table:

'P' means the SCSI Initiator or Target is prohibited from using the specified feature

'R' means the SCSI Initiator or Target is required to support the specified feature.

Table 1 – Class 2 Service Parameters

Class 2 Service Parameter	SCSI Initiator	SCSI Target	Notes
Class validity = 1	R	R	
Service Options			
Intermix Mode = 0	R	R	
Stacked Connect Requests = '00'b	R	R	
Sequential Delivery = 0	A	A	
Dedicated Simplex = 0	R	R	
Camp-On = 0	R	R	
Buffered Class 1 = 0	R	R	
Initiator Control			
Sequence Initiator X_ID reassignment = '00'	R	R	
Initial Responder Process_Associator = '00'b	R	R	
ACK_0/ACK_N capable = '00'b	A	A	
ACK_0/ACK_N capable = '01'b	P	P	
ACK_0/ACK_N capable = '10'b	A	A	
ACK_0/ACK_N capable = '11'b	P	P	
ACK generation assistance = 0	R	R	
Initiator Data compression capable = 0	R	R	
Initiator Data compression History buffer size = '00'b	R	R	
Recipient Control			
ACK_0/ACK_N capable = '00'b	A	A	
ACK_0/ACK_N capable = '01'b	P	P	
ACK_0/ACK_N capable = '10'b	A	A	
ACK_0/ACK_N capable = '11'b	P	P	
X_ID interlock = 0	R	R	

Table 1 – Class 2 Service Parameters

Class 2 Service Parameter	SCSI Initiator	SCSI Target	Notes
Error Policy Supported			
Abort, discard multiple sequences ('00'b)	A	A	
Abort, discard a single sequence ('01'b)	R	R	
Process policy with infinite buffers ('10'b)	P	P	
Discard multiple sequences w/immediate retrans ('11'b)	P	P	
Categories per Sequence = '00'b (one)	R	R	
Data compression capable = 0	R	R	
Data compression History buffer size = '00'b	R	R	
Receive data field size (min)	256	256	
Concurrent Sequences > 0	R	R	
N_Port End-to-end Credit (min)	1	1	
Open Sequences per Exchange > 0	R	R	

4.2 Exchange Management

The following clauses define exchange management requirements and optional behavior for devices operating in accordance with this document.

4.2.1 Initiator Exchange Management

Each exchange shall be timed using the SCSI command timeout value specified for the command. This value shall be the ULP timeout value for the FCP exchange. If a valid FCP_RSP is received before the ULP timer expires the exchange is considered complete after the ACK is sent to the target. If the ULP timer expires before the FCP_RSP is received for the command the initiator shall abort the exchange using the Abort Sequence Protocol and return a proper error indication to the application.

The initiator shall maintain exchange information until the ACK for the FCP_RSP has been sent to the target.

4.2.2 Target Exchange Management

The target shall manage exchanges using the OX_ID/RX_ID pair and shall maintain exchange information until the ACK for the FCP_RSP has been received.

4.3 Sequence Management

The following clauses define sequence management requirements for devices operating in accordance with this document.

Sequence management shall be performed as specified in the FC-PH document for class 2 operation. Each sequence (frame?) shall be timed using the value of E_D_TOV. If an ACK for the sequence is not received within E_D_TOV a sequence timeout has occurred. If a sequence timeout occurs on the first sequence of the exchange (i.e. the FCP_CMD) the initiator shall issue a RES and the Abort Sequence Protocol shall be invoked.

4.3.1 FCP Sequence Delivery Confirmation

Since this document describes operations in a class 2 environment, acknowledgments are used to provide explicit confirmation of sequence delivery. Implicit delivery may be achieved by issuing the RES extended link service and checking the response or the status of the Sequence Initiative bit in the returned exchange status block. For class 2 the sequence recipient shall not consider Sequence Initiative to have been passed until the sequence which passes the Sequence Initiative has completed successfully and the ACK has been transmitted with the Sequence Initiative bit = 1.

5 Error Detection and Recovery

This clause describes the error actions to be taken by the SCSI Initiator and SCSI Target upon detection of an error condition.

5.1 Error Detection by SCSI Initiator

The Abort Sequence Protocol shall be invoked by a SCSI Initiator when any of the following conditions occur:

- a) A Sequence error is detected.
- b) ULP_TOV has expired and an Exchange has not completed.

5.2 Error Detection by SCSI Target

- a) A Sequence error is detected.

When a SCSI Target detects a sequence error, it shall discard that Sequence only and perform the Abort Sequence Protocol. If the target holds sequence initiative it may optionally attempt to send a FCP_RSP with SCSI status and sense information.

5.3 Abort Sequence Protocol

The Abort Sequence Protocol is primarily used in this document to abort a single sequence within an exchange. The unit of error recovery for this document is a sequence not an exchange. A SCSI Initiator may also abort an exchange using ABTS with the LS bit = 1.

A SCSI Initiator and Target may transmit ABTS. ABTS may be transmitted even if Sequence Initiative is not held or there is no end-to-end credit available. Following the transmission of ABTS the port shall consider the status of the exchange indeterminate and shall not deliver any sequences or notification of delivery to the application until the BA_ACC is received, processed, and recovery (if any) is performed. The sender of the ABTS shall reset the E_D_TOV and R_A_TOV timers when the ABTS is transmitted.

Following receipt of the BA_ACC in response to an ABTS, and after R_A_TOV_{SEQ_QUAL} has elapsed, the ABTS Initiator shall transmit RRQ.

If a proper BA_ACC, BA_RJT, LOGO, or PRLO is not received from the recipient within E_D_TOV, second level error recovery as described below shall be performed.

5.3.1 ABTS Recipient Behavior

Upon transmission of any of the above responses, the recipient may reclaim any resources associated with the designated sequence after R_A_TOV_{SEQ_QUAL} has elapsed or a Reinstatement Recover Qualifier (RRQ) extended link service request has been received.

SCSI Targets shall qualify ABTS based upon OX_ID and RX_ID, since the RX_ID is not guaranteed to be known by a SCSI Initiator.

5.3.2 Reinstatement Recovery Qualifier (RRQ)

The RRQ link service is used to notify the destination port that the Recovery_Qualifier is available for reuse. A separate exchange shall be used to reinstate the Recovery_Qualifier. Resources associated with the OX_ID and RX_ID specified shall be released following reception of the ACC at the initiator and following transmission of the ACC at the recipient.

The recovery qualifier is a data structure provided by the recipient to the initiator in the BA_ACC to the ABTS. It describes a completely qualified exchange and sequence and a range of frames

The format of the RRQ is shown in figure 1.

Figure 1 – Reinstatement Recovery Qualifier

	Field	Content
Frame Header	OX_ID	Identifier of a new exchange
	RX_ID	hex 'FFFF'
Payload	Originator S_ID	Source_ID of the Initiator
	OX_ID	OX_ID of XCHG that was previously aborted with ABTS
	RX_ID	RX_ID of XCHG that was previously aborted with ABTS


Following successful completion of the RRQ, the recipient shall respond with ACC.

5.3.3 Second-level error recovery

If a response to the ABTS is not received within E_D_TOV the Initiator shall send the ABTS again. This procedure shall continue until the ULP timer expires.


5.3.4 Error Detection and Recovery Examples

Figure 2 – Lost FCP_CMND




If a FCP_CMND is sent by the initiator and not received at the target the timer will expire and the initiator sends a RES. The target does not know of the exchange and returns a LS_RJT with an invalid OX_ID-RX_ID combination reason code. The initiator knows the command was not received so it performs the Abort Sequence Protocol and reissues the command.

Figure 3 – Lost ACK for FCP_CMND




The FCP_CMND is sent by the initiator and received at the target. The target sends the ACK but it does not arrive at the initiator. The timer will expire and the initiator sends a RES. The target knows of the exchange and returns a ACC. The initiator knows the command was received and performs the Abort Sequence Protocol and maintains the exchange.

Figure 4 – Loss of FCP_XFER_RDY/FCP_RSP (or lost ACK for FCP_XFER_RDY/FCP_RSP)




The FCP_XFER_RDY is sent by the target and not received at the initiator. The target does not see an ACK so the timer expires and it performs the Abort Sequence Protocol. After the ACC for the RRQ is received the target sends the FCP_XFER_RDY again.

Figure 5 – Loss of FCP_DATA (Write)



The FCP_DATA sequence is sent by the initiator but all the data is not received at the target so no ACK will be sent. The timer will expire and the initiator performs Abort Sequence Protocol. After the ACC for the RRQ is received the FCP_DATA sequence is retransmitted using a new sequence ID. (Note: the Tachyon currently requires the FCP_XFER_RDY to be resent also).

Figure 6 – Loss of FCP_DATA (Read)



The FCP_DATA sequence is sent by the target but all the data is not received at the initiator so no ACK will be sent. The timer will expire and the target performs Abort Sequence Protocol. After the ACC for the RRQ is received the FCP_DATA sequence is retransmitted using a new sequence ID.

5.4 Notes

- a) Should work with any ACK model (ACK 0 is preferred).
- b) Need an indication that an exchange was aborted in LS_RJT response (versus lost ACK for FCP_RSP).

6 SCSI Stream Devices

This clause describes commands and features applicable to SCSI stream devices with usage defined by this profile.

6.1 Applicable Classes of Service

SCSI-3 stream devices conforming to this document shall use Class 2 service with parameters as described in Table 3.

6.2 Asynchronous Event Notification (AEN)

The use of AEN by stream devices conforming to this profile is prohibited.

6.3 Command Linking

Command Linking is allowed by all stream devices conforming to this profile. Support for command linking is identified in the Inquiry data. The Flag bit of the CDB shall be set to zero.

6.4 Sequential device commands

Command and features within commands, that are not listed are optional. Interoperability between SCSI Initiators and SCSI Targets is not guaranteed if optional commands or features are used. If support of a field is listed as Required without specifying a value for that field, it is assumed that all possible values of the bit or field shall be supported in accordance with the appropriate American National Standard. Any unlisted commands/features/settings are implicitly Invokeable by the SCSI Initiator, and Allowed by the SCSI Target.

Table 2 – SCSI Tape Device Commands

Feature	Initiator	Target	Doc	Notes
ERASE	I	R		
IMMED	I	R	SSC	
LONG	I	R		
SHORT	I	A		
FORMAT MEDIUM	I	A	SSC	
INQUIRY	I	R		
Standard INQUIRY data (bytes 0-35)	I	R		
EVPD=1	I	R		
Vital Product Data page codes:				
hex'00' (Supported vital product pages)	I	R	SPC	
hex'80' (Unit serial number page)	I	R		
hex'81' (Implemented operations definition page)	I	A		
hex'82' (ASCII implemented operations definition page)	I	A		
hex'83' (Device identification page)	I	A		
LOAD UNLOAD	I	R		
EOT = 1	I	A	SSC	
Immed = 1	I	R		
Load = 1	I	A		
Load = 0	I	R		
Reten = 1	I	A		

Table 2 – SCSI Tape Device Commands

Feature	Initiator	Target	Doc	Notes
LOCATE	I	R		
BT = 0	I	R		
BT = 1	I	A		
CP = 0	I	R	SSC	
CP = 1	I	A		
Immed=0	I	R		
Immed=1	I	R		
LOG SELECT	I	A		
SP = 0	I	R		
SP = 1	I	A		
PCR = 1 & PLL = 0	I	R		
PC = '00'b	I	A	SPC	
PC = '01'b	I	R		
PC = '10'b	I	A		
PC = '11'b	I	A		
LOG SENSE	I	R		
SP = 0	I	R		
SP = 1	I	A		
PPC = 0	I	R		
PPC = 1	I	A		
PC = '00'b	I	A		
PC = '01'b	I	R	SPC	
PC = '10'b	I	A		
PC = '11'b	I	A		
Page Select 00	I	R		
Page Select 02	I	R		
Page Select 03	I	R		
Page Select 06	I	R		
Page Select 0C	I	R		
MODE SENSE/MODE SELECT (6 and 10)	I	R	SPC	
DBD = 0	I	R		
DBD = 1	I	R		
PC = '00'b	I	A		
PC = '01'b	I	R		
PC = '10'b	I	A		
PC = '11'b	I	A		
Page 01	I	R	SPC	
Page 02	I	R	SPC	
Page 0A	I	R	SSC	
Page 10	I	R	SSC	
Pages 11-14	I	A	SSC	
MOVE MEDIUM	I	A	SMC	1
PERSISTENT RESERVE IN	I	R	SPC	
PERSISTENT RESERVE OUT	I	R	SPC	
PREVENT/ALLOW MEDIUM REMOVAL	I	A	SSC	

Table 2 – SCSI Tape Device Commands

Feature	Initiator	Target	Doc	Notes
READ	I	R		
Fixed = 1	I	R		
Fixed = 0 (Required if Read Block Limits values differ)	I	A	SSC	
SILI = 0	I	R		
SILI = 1 (Required if Read Block Limits values differ)	I	A		
READ BLOCK LIMITS	I	R	SSC	
READ BUFFER	I	A	SPC	
READ ELEMENT STATUS	I	A		1
Vol Tag = 0	I	R	SMC	1
Vol Tag = 1	I	A		1
READ POSITION	I	R		
BT = 0	I	R	SSC	
BT = 1	I	A		
READ REVERSE	I	A		
Fixed = 1	I	R		
Fixed = 0 (Required if Read Block Limits values differ)	I	A		
SILI = 0	I	R	SSC	
SILI = 1 (Required if Read Block Limits values differ)	I	A		
Byte Ord = 0	I	A		
Byte Ord = 1	I	A		
RECEIVE DIAGNOSTIC RESULTS	I	A	SPC	
RECOVER BUFFERED DATA	I	R		
Fixed = 1	I	R		
Fixed = 0 (Required if Read Block Limits values differ)	I	A		
SILI = 0	I	R	SSC	
SILI = 1 (Required if Read Block Limits values differ)	I	A		
FIFO	I	A		
LIFO	I	R		
RELEASE(6)	I	A		
3rd Party = 0	I	R	SPC	
3rd Party = 1	P	P		
RELEASE(10)	I	R		
3rd Party = 0	I	R	SPC	
3rd Party = 1	I	A		
REPORT DENSITY SUPPORT	I	R	SSC	
REPORT LUN	I	A	SPC	
RESERVE(6)	I	A		
3rd Party = 0	I	R	SPC	
3rd Party = 1	P	P		
RESERVE(10)	I	R		
3rd Party = 0	I	R	SPC	
3rd Party = 1	I	A		
REWIND	I	R		
Immed = 0	I	R	SSC	
Immed = 1	I	R		

Table 2 – SCSI Tape Device Commands

Feature	Initiator	Target	Doc	Notes
SEND DIAGNOSTIC	I	R		
ST = 1	I	R	SPC	
ST = 0	I	A		
SPACE	I	R		
Forward	I	R	SSC	
Reverse	I	R		
Code = '000'b (Blocks)	I	R		
Code = '001'b (Filemarks)	I	R		
Code = '010'b (Sequential Filemarks)	I	A		
Code = '011'b (End-of-Data)	I	A		
Code = '100'b (Setmarks)	I	A		
Code = '101'b (Sequential Setmarks)	I	A		
TEST UNIT READY	I	R	SPC	
VERIFY	I	A		
Fixed = 1	I	R		
Fixed = 0 (Required if Read Block Limits values differ)	I	A		
BytCmp = 0	I	R	SSC	
BytCmp = 1	I	A		
Immed = 0	I	R		
Immed = 1 & BytCmp = 0	I	R		
Immed = 1 & BytCmp = 1	I	A		
WRITE	I	R		
Fixed = 1	I	R	SSC	
Fixed = 0 (Required if Read Block Limits values differ)	I	A		
WRITE BUFFER	I	A		
Mode = '110'b (Download ucode w/ offsets)	I	R	SPC	
Mode = '111'b (Download ucode w/ offsets & save)	I	R		
WRITE FILEMARK	I	R		
Count = 0	I	R		
Count > 0	I	R	SSC	
Immed = 0	I	R		
Immed = 1	I	R		
NOTES:				
1 Medium changer commands and features are disabled when Inquiry data indicates MCHGR = '0'b and enabled when the MCHGR = '1'b				