

# **FC-AL DIRECT ATTACH DISK PROFILE (Private Loop)**

## **Version 1.3**

Printed: January 27, 1995

### **Contributing Entities:**

**Adaptec  
AT&T-GIS  
Conner  
Cray Research  
EMC  
Emulex  
Hewlett-Packard  
IBM RISC Systems/6000  
NCR Microelectronics  
QLogic  
Seagate  
Storagetek  
Sun Microsystems  
Western Digital**

### **Abstract:**

This Profile selects and restricts options from the Fibre Channel Physical, Fibre Channel Protocol for SCSI, Fibre Channel Arbitrated Loop, and Small Computer Systems Interface standards, such that any device complying with the profile should interoperate according to the interoperability rules specified by the profiles.

This Profile was created by an Ad Hoc group whose members include individuals from the companies listed above. It is not endorsed by FCA or FCSI as of this writing.

Kurt Chan  
M/S 5601  
Hewlett-Packard  
8000 Foothills Blvd  
Roseville, CA 95747-5601

kc@core.rose.hp.com  
916-785-5621 voice  
916-785-2875 fax

This page intentionally left blank

## TABLE OF CONTENTS

1	Introduction	6
2	Scope	6
3	References	6
4	Definitions and Conventions	6
	4.1 Applicability and Use of Profiles	6
	4.2 Implementation of Prohibited Features	7
	4.3 Private versus Public Loop Behaviors	8
5	FC-2 Feature Sets	9
	5.1 General Features	9
	5.2 Extended Link Services	11
	5.3 Basic Link Services	12
6	SCSI Commands and Status	13
	6.1 General Commands	13
	6.2 MODE SELECT/SENSE (10) Parameters (Direct Access)	13
	6.3 SCSI Status	13
7	FC-AL Feature Set	14
8	Initialization	15
	8.1 Initializing LIP	15
	8.2 Selective Hard Reset LIP	15
	8.3 Failure to Obtain Hard-Assigned, Previously-Assigned, or Fabric-Assigned Address	15
	8.4 Failure to Acquire Loop	15
	8.5 Loop Failures	15
	8.6 Target Discovery	16
	8.7 Clearing Effects of ULP, FCP, FC-PH, and FC-AL Actions	17
9	SCSI-FCP Feature Set	18
	9.1 Process Login Parameters	18
	9.2 FCP Information Units	19
	9.2.1 Initiator to Target IUs	19
	9.2.2 Target to Initiator IUs	19
	9.3 Task Management Flags and Information Units	20
	9.4 FCP Task Attributes	21
	9.5 FCP_RSP Payload	21
	9.6 Other FCP Features	22
10	Error Detection	22
11	Error Recovery	22
	11.1 ABTS Protocol	22
	11.1.1 ABTS Frame	23
	11.1.2 BA_ACC frame	23
	11.1.3 BA_RJT frame	23
	11.2 SCSI Initiator Behavior	23
	11.3 SCSI Target Behavior	24
	11.4 SEQ_ID Reuse	24
	11.5 Auto Contingent Allegiance (ACA)	25
	11.5.1 NACA=0 Behavior	25
	11.5.2 NACA=1 Behavior	25
	11.6 Second Level Error Recovery	25
	11.7 Task Management and Multiple-Initiator Targets	25

12	FC-PH 4.3 Deltas	27
13	Drive Connector Requirements	27
13.1	Single Connector Attachment	27
13.1.1	Connector Definition	27
13.1.2	Pinout of Internal Connector	28
13.1.3	Fibre Channel Attachment Design Considerations	29
13.1.4	Fibre Channel Attachment Signal Definitions	29
13.1.4.1	Power	29
13.1.4.2	Fault LED Out	29
13.1.4.3	-Write Protect	29
13.1.4.4	Reserved Signals	29
13.1.4.5	Spindle Sync	30
13.1.4.6	Active LED Out	30
13.1.4.7	PBC Interlock CH1 - CH2	30
13.1.4.8	Motor Start Controls	32
13.1.4.9	Fibre Channel Signals	32
13.1.4.10	SEL_6->SEL_0 ID Lines	32
13.1.5	Backpanel Connector	33
13.1.6	Drive and Backpanel Connector Properties	33
13.2	Direct Attach Cable Specifications	33
13.3	Direct Attach Cable Connectors	34
14	Direct Attach Electrical Requirements	34
14.1	Direct Attach Transmitter Specifications	34
14.2	Direct Attach Receiver Specifications	36
14.3	Equalizer Specifications	37
15	Dual Ports	37
16	External Electrical and Mechanical Requirements	38
Annex A:	Pending ANSI Standards Changes	39
A.1	FC-PH/EP	39
A.1.1	Class 3 Streamed and Open Sequence Definitions	39
A.1.2	Fabric Login Required (23.2, p171)	40
A.2	FCP	40
A.2.1	Task Management and FCP_RSP	40
A.3	SAM	40
A.3.1	ACA ACTIVE FAULTED INITIATOR status	40
Annex B:	Nonzero Login_BB_Credit	42
B.1	Terminology	42
B.2	OPN, CLS, and BB_Flow Control Rules	42
Annex C :	Changes from Previous Revisions	43
C.1	Version 1.20 to Version 1.30	43

## List of Tables

Table 1. Fibre Channel Connector Pin Out	28
Table 2. Output Characteristics of LED Driver Signal	30
Table 3. Definition of Motor Start Controls	32
Table 4. Electronic Requirements for Inputs	32
Table 5. Direct Attach Twinax Electrical Performance (Worst-Case)	33
Table 6. Direct Attach Twinax Electrical Characteristics (Worst-Case)	34
Table 7. FC-AL 1.0625Ghz Transmitter Specifications	35
Table 8. FC-AL Receiver Specifications	37

## List of Figures

Figure 1 - Public and Private Loop Device Coexistence	8
Figure 2 - Class 3 FCP Protocol and Information Units	19
Figure 3 - Target SEQ_ID Reuse on Reads	24
Figure 4 - Port Bypass Circuit Example	31
Figure 5 - Transmitter Eye Diagram	36
Figure 6 - Receiver Eye Diagram	36
Figure 7 - Dual Port Implementation Example	38

## 1 Introduction

This Profile specifies Fibre Channel and SCSI-3 options to ensure interoperability between SCSI Initiators and Disk Targets on a private Loop (no Fibre Channel fabric attached). Future extensions or revisions of this Profile may cover non-disk devices and Public Loop topologies. In FCP-compliant implementations, each SCSI I/O operation is performed through specific Fibre Channel Information Units transferred as Sequences within a single Fibre Channel Exchange created for each I/O operation. These Sequences include a Command Sequence, zero or more Data Transfer Sequences, and a Response Sequence. On Write operations, a Transfer Ready Sequence also precedes each Data Transfer Sequence. Multiple concurrent I/O operations are supported through the use of multiple concurrent open Exchanges.

## 2 Scope

This Profile specifies settings of implementation options required for interoperability with SCSI Initiators and SCSI Disk Drives using FCP connected directly to a common Private Arbitrated Loop. The internal characteristics of those implementations are outside the scope of this Profile. This Profile incorporates features from the standards described below.

## 3 References

- [1] Fibre Channel - Physical and Signalling Interface (FC-PH), X3T9.3/Project 755D/Rev 4.3
- [2] Fibre Channel Protocol for SCSI (FCP), Revision 10, X3T9.2-993D
- [3] Fibre Channel Arbitrated Loop, X3T11/Project 960D/Rev 4.4
- [4] Small Computer Systems Interface - 2 (SCSI\_2), X3.131-1994
- [5] SCSI-3 Architecture Model (SAM), X3T9.2/994D/Rev 16
- [6] SCSI-3 Primary Commands, X3T10/995D/Rev 4
- [7] SCSI-3 Block Commands, X3T10/996D/Rev 0
- [8] Fibre Channel - Extended Physical and Signalling Interface (FC-EP), X3T11

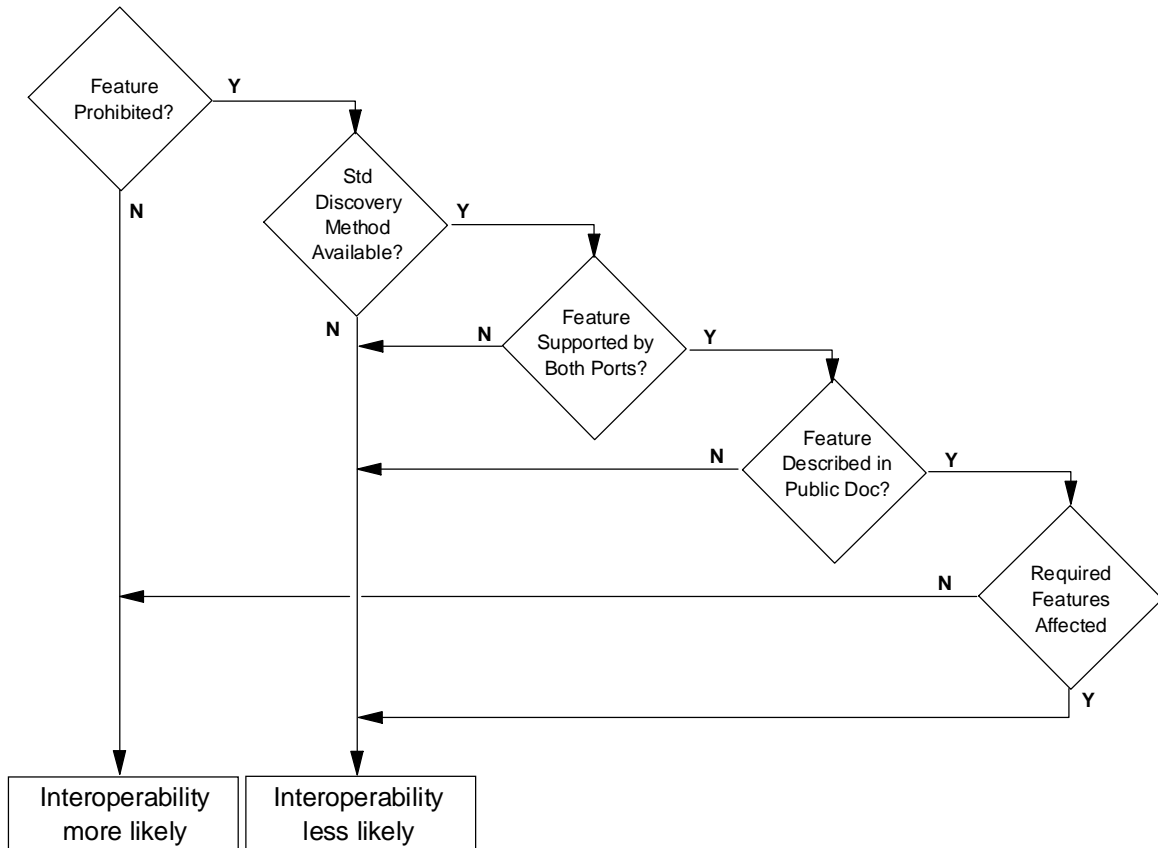
## 4 Definitions and Conventions

### 4.1 Applicability and Use of Profiles

This Profile specifies which FC and SCSI features are required and prohibited in order to ensure interoperability between compliant implementations. Where a feature is not mentioned, it has been determined to have no effect on interoperability. The base structure of this document is taken from the FCSI Profiles.

If a feature is Prohibited, it means that it may not be used between compliant implementations. It does not mean that an implementation may not use that feature to communicate with non-compliant implementations. Interoperability is not guaranteed if Prohibited features are used, or if Required features are not used. This Profile does not prohibit the implementation of features, only their use between compliant implementations. The only features required or prohibited are those which have been determined to result in non-interoperability if implemented differently.

## 4.2 Implementation of Prohibited Features



Features in this profile are prohibited for one of two reasons:

- 1) Simplification: implementing the feature is perceived to increase cost and/or complexity of the implementation, architecture, or documentation. This includes the case where a feature could be made optional but was not, due to lack of a discovery method for that feature, or an incomplete definition in ANSI standards coupled with a lack of contribution outside ANSI standards to make the feature interoperable.
- 2) Interoperability: if one product implements the feature and the other does not, they will not interoperate.

A feature which is Prohibited for reason (1) above may still be used between implementations, as long as:

- a) there is an ANSI standard discovery mechanism defined which all implementations can use to test for use of the optional feature. The ability to discover a capability (or lack thereof) allows an implementation to discover early in the initialization process whether or not it can interoperate with another implementation.
- b) the feature is supported by both implementations
- c) the feature is described in great enough detail in a publicly available document (ANSI Standard, ANSI Technical Report, Complementary Profile, Annex, etc) so that interoperability can be ensured.
- d) use of the feature does not affect the behavior of Required features or interoperability with devices which do not implement the feature.

If all of the above are met, the feature is termed a "Complementary Feature" to the Required feature sets in this profile.

An example of such a feature is nonzero Login\_BB\_Credit. Even though this profile requires Login\_BB\_Credit = 0, Annex B describes how to discover and interoperate with ports which advertise nonzero Login\_BB\_Credit without compromising interoperability with devices which conform to this profile.

### 4.3 Private versus Public Loop Behaviors

The definition of Public and Private Loop behavior is as follows:

	Private Loop Behavior	Public Loop Behavior
Upper 16 bits of native address identifier = 0000h	Required	Prohibited
Opens AL_PA=00	Prohibited	Required (for FLOGI)
FL_Port may Open device	No	Yes

A device which **only** exhibits Private Loop behavior is called a Private Loop Device. A device which exhibits Public Loop behavior is called a Public Loop Device, even though it may communicate with Private Loop Devices. For example, a Public Loop Device below may be an NFS server which communicates with NFS Clients residing directly on the Fabric using IP, and with local SCSI Targets on the same loop using FCP if it follows the rules described in this Profile.

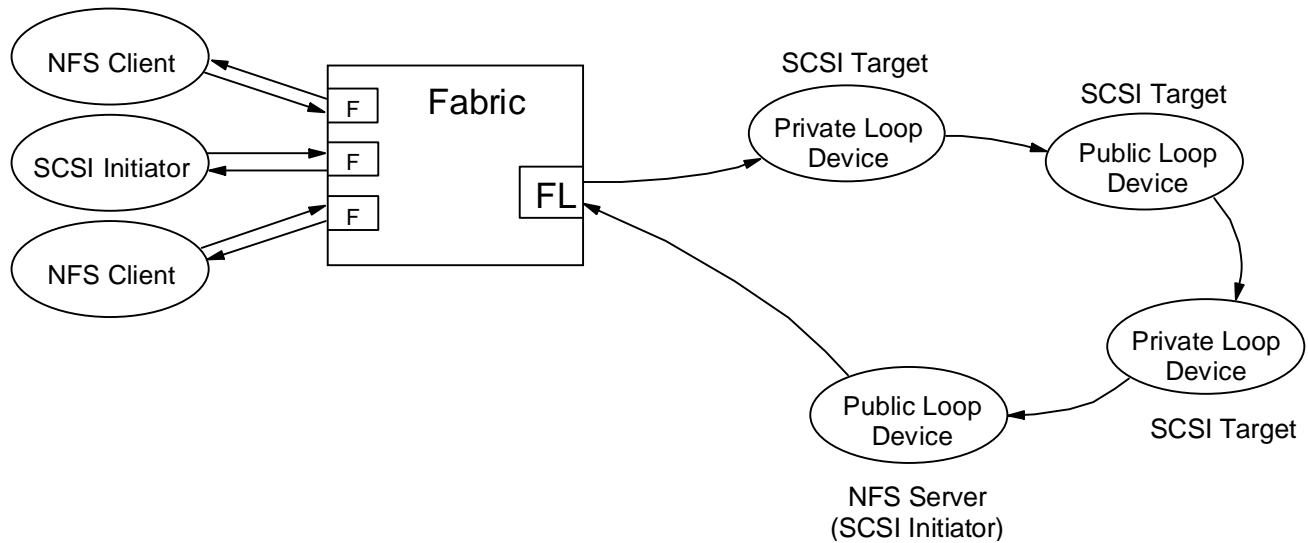


Figure 1 - Public and Private Loop Device Coexistence

The AL\_PA of an NL\_Port corresponds to the lower 8 bits of it's 24-bit native address identifier (with the exception of an FL\_Port or F/NL\_Port which respond to AL\_PA=00). No two ports on the same loop shall share the same AL\_PA, even though the upper 16 bits of their address identifiers may differ.



## 5 FC-2 Feature Sets

The following table lists features described in the Fibre Channel Physical and Signaling Interface standard. This table indicates whether the feature is required, prohibited, or optional for compliance with this specification. Features which are not listed do not affect interoperability of Private Loop devices.

### 5.1 General Features

Reserved FC-PH fields are not required to be checked for zeroes. Validity bits also remove any requirement to check the field for zeroes (e.g., bit 3 of FC-PH F\_CTL = Relative Offset present. If zero, implementations need not check the parameter field in word 5 of the frame header for zeroes).

- R: Use of this feature is Required  
P: Use of this feature is Prohibited  
A: Use of this feature is Allowed (Optional). A corresponding note may recommend behavior.

#### General FC-2 Feature Sets

Feature	SCSI Initiator	SCSI Target	Notes
X_ID reassignment	P	P	
X_ID invalidation	P	P	
Process Error Policy	P	P	
Use of Intermix	P	P	
Class 1	P	P	1
Class 2	P	P	1
Class 3	R	R	
R_CTL Routing Bits			
FC-4 Device_Data frame	R	R	
Extended Link_Data frame	R	R	
Basic Link_Data frame	R	R	
Link_Control frame	R	R	
Lowest FC-PH Version supported	4.3	4.3	
Highest FC-PH Version	4.3	4.3	
Vendor Version Valid = 1	R	R	2
Vendor Code = xx	R	R	2
Continuous Relative Offset	R	R	
Random Relative Offset	P	P	
Relative Offset by Category	R	R	
Generation of nonzero Continue Sequence Condition values	P	P	
Ignore nonzero Continue Sequence values	R	R	

X_ID Interlock	P	P	4
Sequence Chaining (C_S bit in F_CTL)	P	P	
More than one Information Category per Sequence	P	P	
Continuously Increasing SEQ_CNT	A	A	
Max frame size used on all but last frame of Sequence	R	R	5
Optional Headers (all)	P	P	
Node Name Format			
IEEE	A	P	
IEEE Extended	A	R	3
N_Port Name Format			
IEEE	A	P	
IEEE Extended	A	R	3

Notes:

1. Private Loop Initiators are prohibited from initiating Class 1 or 2 frames. PDISC may be issued in any class of service to discover login parameters. Therefore, there is no need to allow class 1 or 2 operation in Targets which only support private loop operation.
2. Devices which support this profile shall indicate so through the vendor version level bits (see FCSI method proposed at 12/94 X3T11 meeting)
3. On a dual-port implementation, both ports share the same lower 48 bits of the IEEE Extended address.
4. X\_ID interlock only applies to classes 1 and 2. If RX\_ID <> FFFF is used by an Exchange Responder, it shall return the RX\_ID value on the first XFER\_RDY sequence (for write operations) or the first Data sequence (for read operations).
5. Only the last frame of a Sequence may be less than "full" size, where "full" is defined by the buffer-to-buffer receive data field size in the N\_Port common service parameters (word 1, bits 15-0).

## 5.2 Extended Link Services

P: Device is prohibited from transmitting (receiving device not required to accept)

R: Device is required to have the ability to Transmit (receiving device is required to accept)

Feature	SCSI Initiator	SCSI Target	Notes
ABTX	P	P	
ADVC	P	P	
ECHO	P	P	
ESTC	P	P	
ESTS	P	P	
FLOGI	P	P	1
LOGO	R	R	6
PDISC	R	P	5
PLOGI	R	P	5
PRLI	R	P	5,7
PRLI Common Service Parameters	P	P	
Multiple Service Parameter pages per PRLI request	P	P	
Accept Response code = "Command executed"	R	R	
ACC contains only those pages specified in PRLI	R	R	
PRLO	R	R	3
RCS	P	P	
RES	P	P	
RLS	R	R	2
RRQ	P	P	
RSI	P	P	
RSS	P	P	
RTV	P	P	4
TEST	P	P	

### Notes:

1. Private Loop devices are prohibited from performing FLOGI
2. Not all fields in the LESB may be completely supported. See vendor-specific documentation.
3. PRLO is the required response if a frame is received from a device which has not performed Process Login.
4. A SCSI Initiator shall use the largest EDTOV presented in the ACC payloads to PLOGI during Target Login. A Loop Target does not initiate action based on Sequence Timeouts. On a private loop, R\_A\_TOV=2 x E\_D\_TOV.
5. See the section on Target Discovery for how these ELSs are used.
6. LOGO is the required response if a frame is received from a device which has not performed N\_Port Login.
7. See FCP for a description of the PRLI Extended Link Service

### 5.3 Basic Link Services

P: Device is prohibited from transmitting (receiving device not required to accept)

R: Device is required to have the ability to Transmit (receiving device is required to accept)

	SCSI Initiator	SCSI Target	Notes
NOP	P	P	
RMC	P	P	
ABTS	R	P	1

1. See clause on Error Recovery.

## 6 SCSI Commands and Status

### 6.1 General Commands

The use of Asynchronous Event Notification is Prohibited.

The following SCSI commands are required by this profile, Commands or features within commands which are not listed are optional.

- FORMAT UNIT (F/C/DLF = 0/0/000)
- INQUIRY (EVPD Serial# page, standard page bytes 0-35)
- MODE SENSE/SELECT(10)
- READ(10)
- READ DEFECT DATA
- RELEASE/RESERVE (Unit only)
- REQUEST SENSE
- SEND DIAGNOSTIC (selftest)
- READ CAPACITY
- START/STOP UNIT
- TEST UNIT READY
- WRITE(10)
- WRITE BUFFER (download microcode, download/save)

### 6.2 MODE SELECT/SENSE (10) Parameters (Direct Access)

"R" means that a Target must be able to accept this parameter in a MODE SELECT command, and must be able to return a value for it in a MODE SENSE command. All unlisted features are optional.

Parameter	Target	Notes
Number of Block Descriptors	1	1
Device Specific Parameter (WP bit)	R	
Control Mode Page (EECA bit = 0)	R	
Disconnect/Reconnect Page		
Buffer Full/Empty Ratio	R	
Maximum Burst Size	R	2

Notes:

1. The Block Length in the Block Descriptor (actual size of a logical block) shall be a integer multiple of 2 bytes.
2. The Maximum Burst Size corresponds to the maximum Sequence Length of the Target, in increments of 512 bytes.  
For class 3, this represents the maximum amount of data the Target is allowed to transmit in a single Loop tenancy.

### 6.3 SCSI Status

Required STATUS values are:

- ACA ACTIVE
- ACA ACTIVE FAULTED INITIATOR (if NACA=1 is used - see clause on ACA and Annex A)
- RESERVATION CONFLICT
- BUSY
- GOOD
- CHECK CONDITION
- TASK SET FULL

## 7 FC-AL Feature Set

R: Use of this feature is Required

P: Use of this feature is Prohibited

A: Use of this feature is Allowed (Optional).

Feature	SCSI Initiator	SCSI Target	Notes
Open Full Duplex	R	R	3
Open Half Duplex	P	P	2
Unfairness	A	A	
Transfer mode	A	P	
LILP/LIRP	P	P	
Multicast/Selective Replicate	P	P	
Alternate BB_Credit model	R	R	1
BB_Credit = 0 on N_Port login	R	R	2

- 1) Alternate BB\_Credit management is mandatory in FC-AL
- 2) Interoperability can be achieved with devices which advertise nonzero BB\_Credit during login if half duplex OPNs are used. See Annex B for details.
- 3) When Login\_BB\_Credit = 0, the OPN Recipient shall either transmit the number of R\_RDY's back to the OPN Initiator equal to the number of available receive buffers, or if the number of available receive buffers is zero, it shall respond with an immediate CLS rather than keep the circuit open until a buffer becomes available.

## **8 Initialization**

### **8.1 Initializing LIP**

LIP(F7,F7) is the initializing LIP. It shall have no effect on existing tasks except that the I/O in progress may be disrupted, resulting in Sequence or ULP timeout. If a device is unable to win arbitration within E\_D\_TOV, it may issue this LIP.

### **8.2 Selective Hard Reset LIP**

LIP(AL\_PD,AL\_PS) is a power-on reset of the device at AL\_PD. All other devices shall treat this as LIP(F7,F7). Targets are prohibited from issuing this LIP.

### **8.3 Failure to Obtain Hard-Assigned, Previously-Assigned, or Fabric-Assigned Address**

A Hard-Assigned AL\_PA is one which is preconfigured via address switches or nonvolatile memory. See Annex K of FC-AL for which switch settings correspond to which AL\_PAs. Address reconfiguration may occur upon receipt of a LIP or upon power-up.

If a hard-assigned AL\_PA cannot be obtained, this indicates a configuration error. Devices shall:

- 1) Accept a new soft address. If no soft address is available, go to non-participating mode.
- 2) Implicitly log out all ports
- 3) Accept new PLOGIs (Targets only)
- 4) Accept any PRLI (Targets only) with a reason code 0101 in word 1, bits 11:8 of the PRLI accept response code ("The target image has a predefined configuration which precludes establishing this image pair. The PRLI request shall not be retried.").

If a previously-assigned AL\_PA cannot be obtained when soft addresses are being used, this indicates a configuration change from the previous configuration, not necessarily a configuration error. Devices shall only perform steps 1-3 above.

### **8.4 Failure to Acquire Loop**

If performance has deteriorated and it is suspected that another port on the loop may not be observing fairness (due to a malfunction or otherwise), then LIP(F7,AL\_PS) may be issued. This does not affect any tasks other than the one in progress on the loop.

### **8.5 Loop Failures**

A Loop Failure is defined as detection of Loss of Signal or Loss of Synchronization for longer than R\_T\_TOV. If a Loop Failure occurs and a valid AL\_PA has been obtained, LIP(F8,AL\_PS) shall be issued. If a Loop Failure occurs and no valid AL\_PA has been obtained, LIP(F8,F7) shall be issued.

## 8.6 Target Discovery

Whenever a possibility of a topology change exists, a SCSI Initiator may want to rediscover the new topology. The Target Discovery procedure for a SCSI Initiator is:

```
If port resides on Loop, then for all legal AL_PAs except 00h,
    OPN(AL_PA)
    IF OPN successful, THEN
        PDISC(AL_PA)
        IF new WWN or login parameters changed at AL_PA, THEN
            PLOGI(AL_PA) in class 3
            IF PLOGI successful, THEN
                PRLI(AL_PA)
                IF PRLI successful, THEN
                    MODE SENSE(AL_PA)
                    INQUIRY(AL_PA)
                ENDIF
            ENDIF
        ENDIF
    ENDIF
ENDIF
NEXT AL_PA
```

where PDISC is the Port Discovery Extended Link Service described in FC-EP. This procedure is designed to avoid the abnormal termination of all open Exchanges whenever a new device is inserted or removed from the loop, or whenever a device powers on or off.

SCSI Targets are prohibited from performing N\_Port login, while SCSI Initiators are required to be able to perform N\_Port login.

Upon receipt of a LIP, a Target shall suspend execution of any current tasks and shall not allow any pending tasks to become current tasks until PDISC is received from the Initiator of the current or pending task.



## 8.7 Clearing Effects of ULP, FCP, FC-PH, and FC-AL Actions

Y = Object is cleared to its default or power-on value upon completion of specified action. In a dual-port implementation, the object is cleared regardless of which port the action occurs on.

<b>Object</b>	<b>Action</b>	Reset LIP(y,x) <sup>7</sup>	LOGO, PLOGI	ABTS <sup>9</sup>	PRLI, PRLO	Target Reset	Clear Task Set <sup>9</sup>	Abort Task Set <sup>9</sup>
PLOGI parms		Y <sup>4</sup>	Y	N	N	N	N	N
Open Sequences		Y <sup>4</sup>	Y <sup>2</sup>	Y <sup>2</sup>	Y <sup>2</sup>	N <sup>3</sup>	N <sup>3</sup>	N <sup>3</sup>
EE_Credit_CNT		Y	Y	N	N	N	N	N
BB_Credit_CNT		Y	Y	N	N	N	N	N
AL_PA		N <sup>1</sup>	N	N	N	N	N	N
PRLI parms		Y <sup>4</sup>	Y <sup>2</sup>	N	Y	N	N	N
Open Tasks		Y <sup>4,5</sup>	Y <sup>2,5</sup>	Y <sup>2,6</sup>	Y <sup>2,5</sup>	Y <sup>4,5</sup>	Y <sup>4,5</sup>	Y <sup>2,5</sup>
Mode page parms = saved		Y <sup>4</sup>	Y <sup>2</sup>	N	Y <sup>2</sup>	Y <sup>4</sup>	N	N
Pre-existing UA Condition		Y	Y <sup>8</sup>	N	Y <sup>8</sup>	Y	N	N
Pre-existing ACA Condition		Y <sup>4</sup>	Y <sup>8</sup>	N	Y <sup>8</sup>	Y	N	N

Notes:

1. Unless AL\_PA's have changed (i.e., preferred address has already been taken by another port)
2. For Sequence, Process, or N\_Port Login Initiator only, not for all Initiators.
3. Tasks are cleared internally within the Target, but open FC-PH Sequences must be individually aborted by the SCSI Initiator via ABTS.
4. For all Initiators
5. All open tasks
6. When ACC with L\_S=1 returned
7. This is also known as LIP(AL\_PD,AL\_PS). If the destination of the selective reset LIP matches the address of the receiving device, the receiving device shall observe the behavior described in this column. If the LIP is not addressed to the receiving device, the receiving device shall treat the LIP as it would any other LIP (which does not perform a power-on reset).
8. If the Target has established an ACA or UA condition for an Initiator, and the Initiator logs out (PRLO or LOGO), the ACA/UA condition shall be cleared for that Initiator.
9. For multi-LUN targets, these actions affect only the addressed LUN, not the entire target.

## 9 SCSI-FCP Feature Set

The following Feature Sets are described in the Fibre Channel Protocol for SCSI (FCP). This table indicates whether the feature is required, prohibited, or optional for compliance with this specification.

### 9.1 Process Login Parameters

R: Use of this feature is Required

P: Use of this feature is Prohibited

I: This feature may be invoked by a SCSI Initiator, and therefore must be implemented by a SCSI Target.

Feature	SCSI Initiator	SCSI Target	Notes
Originator Process Associator Valid	P	P	
Responder Process Associator Valid	P	P	
Originator Process Associator	P	P	
Responder Process Associator	P	P	
Initiator Function	R	P	
Target Function	P	R	
Data Overlay Allowed	P	P	
Data + Response in same Sequence (Read)	P	P	
Command + Data in same Sequence (Write)	P	P	
Read XFER_RDY Disabled	R	R	
Write XFER_RDY Disabled	P	P	
FCP_RSP to Task Management	R	R	1

1. This is a feature proposed for FCP - see Annex A.

## 9.2 FCP Information Units

All T1, I1, and I4 Sequences shall be single-frame Sequences. T6 and I3 Sequences may be multiple-frame Sequences.




Figure 2 - Class 3 FCP Protocol and Information Units

### 9.2.1 Initiator to Target IUs

T1 (CMD/Task Mgmt with SI transferred), and T6 (Write Data with SI transferred) are required. All others are prohibited. FCP\_DL in the FCP\_CMND payload shall always be equal to the number of bytes transferred for the command. Transfer Length in the FCP\_CDB shall be expressed in logical blocks. Therefore FCP\_DL shall be equal to Transfer\_Length \* Block\_Size. Targets are not required to check for this consistency.


### 9.2.2 Target to Initiator IUs

I1 (XFR\_RDY on Write), I3 (Read Data), I4 (FCP\_RSP) are required. All others are prohibited. Task Management functions shall be responded to with I4.

On transfers consisting of multiple Write Data Sequences, the DATA\_RO parameter contained in consecutive XFR\_RDY Sequences does not have to be continuously increasing. This provides the logical equivalent of SCSI Modify Data Pointers. This is not the same as the Random Relative Offset N\_Port login Service parameter. Within an Exchange, only the FCP\_DATA frame containing the largest Relative Offset may contain fill bytes.

The DATA\_RO parameter in the XFR\_RDY payload on Write operations shall be a multiple of 8 bytes. The initial Relative Offset of each Read Data Sequence shall also be a multiple of 8 bytes.

### 9.3 Task Management Flags and Information Units



All Initiators shall send Task Management functions using T1. All Targets shall return FCP\_RSP to Task Management functions using I4. RSP\_CODE in the FCP\_RSP\_INFO field shall be as follows:

RSP_CODE	SAM Mapping	Meaning
00	Function Complete	Task management function was performed
04	Function Rejected	Task management function was not performed because the function is unsupported
05	Service Delivery or Target Failure	Task management function was not performed, but the function is supported. For example, if an ACA condition exists for an initiator, the target will return this RSP_CODE to another initiator attempting to perform a Clear Task Set function

#### FCP Task Management Flags

Feature	SCSI Initiator	SCSI Target	Notes
Terminate Task	P	P	
Clear ACA	I	R	
Target Reset	I	R	1
Clear Task Set	I	R	1
Abort Task Set	I	R	

Notes:

1. See the clause on error recovery for Initiator requirements following transmission of these Task Attributes.

## 9.4 FCP Task Attributes

Feature	SCSI Initiator	SCSI Target	Notes
Untagged	P	P	1
Simple Queue Type	R	R	
Ordered Queue Type	I	R	3
Head of Queue Type	I	R	2
Auto Contingent Allegiance Type	R	R	

### Notes:

1. All FCP operations are implicitly tagged with an Exchange ID. Therefore, SCSI Initiators using untagged queuing rules should use the Simple Queue task attribute, and not issue overlapped commands.
2. Head of Queue Type is useful over private loops which can ensure delivery order, but may be less useful across fabrics/topologies which cannot.
3. Ordered Queue Types are allowed to be invoked by SCSI Initiators on private loops, since frame/Sequence/Exchange ordering is preserved.

## 9.5 FCP\_RSP Payload

Feature	SCSI Initiator	SCSI Target	Notes
FCP_SNS_LEN	<=84	<=84	1
FCP_RSP_LEN	<=20	<=20	
FCP_SNS_INFO	R	I	
FCP_RSP_INFO	R	I	2
Length of Additional Sense Bytes in FCP_RSP_INFO	<=66	<=66	1

### Notes:

1. This length restriction ensures the FCP\_RSP payload will be less than or equal to 128 bytes.
2. If present, this field shall contain 4 bytes for FCP error codes, plus up to 16 bytes reserved for future use. The length of the field shall be 4, 8, 12, or 16 bytes.

## 9.6 Other FCP Features

Feature	SCSI Initiator	SCSI Target	Notes
FCP_LUN	R	R	
FCP_LUN (0)	I	R	
INQUIRY of FCP_LUN (0)	I	R	
Process Login	R	R	
Process Logout	P	P	1

1. The single exception to this rule is that PRLO may be used in response to frames from a device which has not performed PRLI as a notification. A PRLI has the effect of implicit PRLO.

## 10 Error Detection

The ABTS protocol (described below in the Error Recovery clause) shall be invoked by a SCSI Initiator when the ULP timer has expired and FCP\_RSP has not been received.

Targets are not required to time Sequences as Sequence Recipients or as Sequence Initiators. Targets are not required to either assign RX\_IDs or detect overlapped commands (duplicate OX\_ID/RX\_ID combinations). SCSI Initiators are required to guarantee X\_ID uniqueness by only reusing OX\_IDs to a Target N\_Port after the Exchange has been terminated (either successfully or unsuccessfully). This is the only restriction on OX\_ID use.

Note: FC-PH, section 21.1.2 on page 133 requires that "the Originator of an Extended Link Service Exchange shall detect an Exchange error following Sequence Initiative transfer if the Reply Sequence is not initiated and received within a timeout interval equal to twice the value of R\_A\_TOV."

## 11 Error Recovery

The clause on Loop Initialization contains other requirements for Loop-specific error recovery.

### 11.1 ABTS Protocol

The ABTS-L\_S protocol is required, which uses ABTS to abort entire Exchanges. The unit of error recovery for this Profile is an Exchange, not a Sequence. The protocol defining retransmission of Exchanges following ABTS is beyond the scope of this document.

Only a SCSI Initiator may initiate FCP Recovery Abort. This Profile does not define the protocol by which multiple Initiators communicate or synchronize shared peripherals.

### 11.1.1 ABTS Frame

SEQ\_ID = SEQ\_ID of an Open Sequence at the SCSI Initiator. If no Sequence is open, then the SEQ\_ID is any SEQ\_ID not currently open between the SCSI Initiator and Target.

SEQ\_CNT = SEQ\_CNT of last frame transmitted in an Open Sequence + 1. If no Sequence is open, then SEQ\_CNT = 0.

OX\_ID = same as that assigned by SCSI Initiator (Exchange Originator) for the command being aborted.

RX\_ID = FFFF if no XFR\_RDY (on Write commands) or Read Data (on Read commands) has been received from Target. Otherwise, RX\_ID = same as that assigned by Target for the I/O being aborted.

F\_CTL Sequence Context = Initiator (even though the port transmitting ABTS may not have Sequence Initiative for the Sequence being aborted).

### 11.1.2 BA\_ACC frame

Only a SCSI Target is required to accept ABTS with BA\_ACC.

Frame Header:

- OX\_ID in frame header = OX\_ID from ABTS frame.
- RX\_ID in frame header = RX\_ID from ABTS frame

Payload:

- SEQ\_ID Validity bit in payload = 00 hex
- SEQ\_ID byte in payload = invalid (don't-care)
- OX\_ID in payload = OX\_ID from ABTS frame
- RX\_ID in payload = RX\_ID from ABTS frame
- Lowest SEQ\_CNT in payload = 0
- Highest SEQ\_CNT = FFFF

F\_CTL:

- L\_S bit set to indicate last frame of Exchange.
- Sequence Context = Recipient

### 11.1.3 BA\_RJT frame

BA\_RJT is transmitted by a Target in response to ABTS only if:

- 1) The Target has assigned an RX\_ID to an OX\_ID in a previous frame, and
- 2) The SCSI Initiator has issued ABTS which contains an unknown OX\_ID/RX\_ID combination

The reason code shall be "Logical Error" with a reason code explanation of "Invalid OX\_ID/RX\_ID combination".

## 11.2 SCSI Initiator Behavior

ABTS may be transmitted even if Sequence Initiative is not held. Following the transmission of ABTS, any Device\_Data frames received on this Exchange shall be discarded until the BA\_ACC with L\_S=1 bit is received from the Target. If a proper BA\_ACC or BA\_RJT is not received from the Target within (2 x R\_A\_TOV), second level error recovery shall be initiated.

Initiators are prohibited from using the RRQ protocol since frames on a private loop cannot be reordered. Sequence Qualifiers are reusable as soon as the ACC is returned to ABTS.

### 11.3 SCSI Target Behavior

When ABTS is received at the Target, it shall return either BA\_ACC or BA\_RJT. If RX\_ID has been assigned, and an invalid OX\_ID/RX\_ID combination has been detected, then BA\_RJT shall be transmitted. Otherwise, all Exchange resources shall be reclaimed and BA\_ACC shall be transmitted with L\_S=1 set in F\_CTL.

If an error is detected by a Target while it has Sequence Initiative, the only permissible recovery action is the transmission of FCP\_RSP with CHECK CONDITION status and an appropriate Sense Key/ASC/ASCQ.

If an error is detected by a Target while it does not have Sequence Initiative, it must wait until it has been given Sequence Initiative until it can return CHECK CONDITION.

If ULP resources in the Target are unavailable when a command is issued, the Target shall return TASK SET FULL status in response to the command. A ULP resource at the Target is a buffer associated with storing the SCSI Command prior to execution, or any buffer required to maintain Command context during execution.

If ULP resources in the Target are unavailable when Login is attempted, the Target may either:

- 1) return LS\_RJT, with reason code "Unable to perform command request" and reason code explanation "insufficient resources to support Login", or
- 2) implicitly log out the "oldest" Initiator in order to accept the new Login. Any new or partial Sequences received from the logged out Initiator shall be discarded, and may be responded to with LOGO.

### 11.4 SEQ\_ID Reuse

For single-frame Sequences,

- a) A device may reuse a SEQ\_ID immediately following transmission of the single-frame Sequence

For multi-frame Sequences,

- b) A device may reuse a SEQ\_ID for which it has Sequence Initiative following confirmation of transfer of Sequence Initiative.
- c) A Target may reuse a SEQ\_ID for a Read Data Sequence immediately following transmission of FCP\_RSP.

Behavior (c) is justified through the following argument (see Figure 3 below):

- All Sequences in FCP are qualified by an X\_ID
- All open X\_IDs between N\_Port pairs must be unique.
- Since Exchange IDs are always checked before an Sequence within the Exchange is delivered to the ULP, SEQ\_IDs are not required to be unique across all Exchanges.
- Targets are prohibited from retransmitting FCP\_RSP, so the previous Exchange is guaranteed to never complete.




Figure 3 - Target SEQ\_ID Reuse on Reads



## 11.5 Auto Contingent Allegiance (ACA)

- During the discovery process, FCP Initiators shall use INQUIRY to discover whether or not Targets support ACA behavior.
- ACA shall be used (NACA=1 in CDB) only when both the sending Initiator and receiving Target support ACA functionality.
- ACA shall not be used (NACA=0 in CDB) when either the sending Initiator or the receiving Target do not support ACA functionality.

*[ed note: the behavior below is still under discussion)*

### 11.5.1 NACA=0 Behavior

If a Target receives a command with NACA=0 and that command fails,

- 1) The Target shall establish a contingent allegiance condition for the faulted initiator
- 2) The Target shall return CHECK CONDITION in the FCP\_RSP along with autosense data.
- 3) Autosense data and the contingent allegiance condition are cleared upon transmission of the FCP\_RSP frame.

### 11.5.2 NACA=1 Behavior

If a Target receives a command with NACA=1 and that command fails, the Target shall:

- 1) Establish an ACA Condition for the faulted initiator.
- 2) Return CHECK CONDITION in the FCP\_RSP along with autosense data.
- 3) Preserve the ACA condition until one of the following:
  - a) Receipt of a CLEAR ACA Task Management function
  - b) Receipt of a TARGET RESET Task Management function
  - c) Receipt of a LIP hard reset
  - d) Powerfail
- 4) Preserve sense data until the ACA condition is cleared. Sense data shall **not** be cleared upon receipt of REQUEST SENSE in order to preserve the integrity of the sense data until confirmation by the faulted Initiator.

While an ACA condition exists, a Target shall return the following status in response to a newly received task:

- a) ACA ACTIVE status if the ACA Condition exists for another Initiator,
- b) ACA ACTIVE FAULTED INITIATOR status if the ACA Condition exists for the Initiator of the new task and the new task is a command other than REQUEST SENSE with the ACA task attribute set.

While an ACA condition exists, a Target shall not return status for any task which was already in the Task Set when the ACA condition was established.

## 11.6 Second Level Error Recovery

If the ABTS protocol fails, the SCSI Initiator shall:

1. Retry the ABTS protocol once.
2. If the retried ABTS protocol fails, the SCSI Initiator shall implicitly log out (FC-PH Logout) the Target. All outstanding Exchanges with that Target are implicitly terminated at the SCSI Initiator.

## 11.7 Task Management and Multiple-Initiator Targets

If a Target Reset or Clear Task Set management function is received by a Target which has multiple Initiators logged in with it, then:

- 1) In class 3, Task Management functions are sent using the I1 Information Unit (FCP\_CMND with Sequence Initiative transferred).

- 2) The Target shall create a Unit Attention Condition for all other Initiators, regardless of whether or not the Target believes a command is outstanding to those Initiators. This is because FCP\_RSP may have been transmitted but not received by the Initiator, or the Initiator may have transmitted a command which has not yet been received by the Target.
- 3) The Target may clear all resources associated with the cleared Exchanges, per SAM.
- 4) The Target shall return FCP\_RSP upon completion of (1) and (2). The payload shall be zeroes. If the Initiator does not receive FCP\_RSP, it may retry the Task Management. If the retry fails, it shall initiate second level error recovery.
- 5) Upon discovery of a power-on Unit Attention Condition, Initiators shall issue ABTS for all commands which are outstanding to that Target. From an Initiator perspective, this refers to all commands for which FCP\_RSP has not been received.
- 6) If the Target receives an ABTS frame for which the referenced S\_ID | OX\_ID | RX\_ID does not exist, it shall treat the ABTS as if it were for an FCP\_CMD which was transmitted but discarded due to CRC error (BA\_ACC with L\_S=1).

## 12 FC-PH 4.3 Deltas

The following clause describes those areas where the profile requires behavior which is prohibited by FC-PH 4.3, or areas where the profile prohibits behavior which is required in FC-PH 4.3. These "deltas" may be incorporated in FC-EP, depending on X3T11 committee consensus.

## 13 Drive Connector Requirements

### 13.1 Single Connector Attachment

The Fibre Channel Single Connector Attachment (FCSCA) system includes signals for a complete Fibre Channel disk drive connection. In addition, +5V and +12V power, spindle synch, motor start control, PBC Interlock, and an LED drive signal are included in the connector definition. The FCSCA connector is designed and placed to allow plugging a drive directly into a backplane. The FCSCA connector provides the necessary electrical connection, but mechanical stability and device retention must be provided by other mechanisms, including mounting brackets, guide rails, clips, or screw attachments.

The connector selected for the FCSCA is the AMP Champ 0.050" Series 1. Various connector options are available to meet the different mounting requirements of the connector to the Fibre Channel drive and the different drive plugging requirements. Other connector manufacturers have compatible connector designs.

The FCSCA connector will allow drive-to-board mating.

The FCSCA is designed principally for the direct plugging of drives into a backplane. Even though a hot plugging capability is being investigated, it is still recommended that power be removed before removing or inserting a drive.

#### 13.1.1 Connector Definition

The drive connector is a vertical or straddle mount plug, part number AMP Champ 5-175473-5 or AMP Champ 557114-4. The connector to which the drive plug mates is the AMP Champ 5-175475-5 vertical receptacle.

The connector technology meets the following requirements:

- 40 signal contacts to provide power and all interface and control signals
- Small connector to fit in the form factor of 1" high 3.5" disk drives, 1.6" high 3.5" disk drives, and 2.5" disk drives.
- Vertical or straddle mount device connector to allow flexibility in drive design
- Shrouded contacts for mechanical protection of mating surfaces
- Polarized housing to prevent incorrect insertion of drive into backplane sockets
- Tolerant alignment guide-in to allow blind mating
- Greater than 1200 mating cycles
- Acceptable insertion and withdrawal forces (90 and 20 grams/contact)
- Electrical properties suitable for the application
- One row of pins is shorter than the other row to facilitate "hot plugging" and allow ground and current-limited power to be plugged first before non-current limited pins make contact.

### 13.1.2 Pinout of Internal Connector

The FCSCA connector pinout accommodates the signals for Dual-Ported Fibre Channel In/Out, +5v power, +12v power, battery power, spindle sync, spin-up control, PBC Interlock control, write protect, Sel I.D., and LED drivers. The pin out is shown below:

**Table 1. Fibre Channel Connector Pin Out**

Connector Contact	Signal Name	Signal Name	Connector Contact
1*	PBC Interlock CH1	12 Volts Charge	21
2*	12 Volts	12V Ground	22
3*	12 Volts	12V Ground	23
4*	12 Volts	+CH1_IN	24
5*	12 Volt Ground	-CH1_IN	25
6*	Battery**	I/O Ground	26
7*	Active LED Out	+CH2_IN	27
8*	Spindle Sync	-CH2_IN	28
9*	RMT_Start	Logic Ground	29
10*	DLYD_Start	+CH1_OUT	30
11*	PBC Interlock CH2	-CH1_OUT	31
12*	SEL_6	Logic Ground	32
13*	SEL_5	+CH2_OUT	33
14*	SEL_4	-CH2_OUT	34
15*	SEL_3	Logic Ground	35
16*	Fault LED Out	SEL_2	36
17*	-Write Protect	SEL_1	37
18*	Reserved NC	SEL_0	38
19*	5 Volts	Reserved NC	39
20*	5 Volts	5 Volts Charge	40

\*= Short Pins

\*\*= Optional

### 13.1.3 Fibre Channel Attachment Design Considerations

No equalizer is present on the disk drive. The equalizer is either on the backpanel or at the bulkhead, for termination of long data lines.

The backpanel applications will use power planes for power distribution.

### 13.1.4 Fibre Channel Attachment Signal Definitions

#### 13.1.4.1 Power

Power is supplied through the FCSCA connector with support for +5 volts and +12 volts. All of the voltage conducting paths utilize shorter contacts on the power handling side of the FCSCA connector to achieve power surge reduction and to facilitate "hot plugging" of the drives. There are longer voltage contacts in the connector to enable charging the capacitors through series resistive paths to limit in-rush currents.

An optional battery pin is provided for the purpose of supplying power to maintain the cache memories on the devices in the event of a DC power loss. On devices equipped for the battery backup of the cache, a DC regulator powers the cache memory and the dynamic circuitry required to refresh the cache memory. The battery voltage must range from 14v maximum to 6 v minimum and be capable of supplying 150 ma (maximum) of current per device.

Four 12 volt signals provide +12 volt power to the device. The current return for the +12 volt power supply is through the 12 volt ground signals. The maximum total current that can be provided to a drive through the 12 volt signal pins is 3 amps. The supply current and return current must be distributed as evenly as possible among the pins. The maximum current typically occurs while the drive motor is starting. Three of the outer pins are shorter to allow capacitive pre-charging through a resistor limited path of the longer +12 Volt Charge contact.

Three 5 volt lines provide logic power to the device. The current return for the +5 volt power supply is through the 5 volt ground signals. The return current must be distributed as evenly as possible among the voltage and ground pins. The maximum total current that can be provided to a drive through the +5 volt signals is 2 amps. The +5 Volt Charge contacts are longer to allow capacitive pre-charging through resistor limited paths before the primary power contacts (shorter) begin to conduct.

The maximum current specified is related to the connector's characteristics. The use of resistor limited paths on the longer contacts of the +12 v and +5 v ( serves to limit in-rush currents as on-board capacitors are charged up. Additional limitations may be associated with the power supply's current budgets and with the system's power dissipation budget. Those limitations are not controlled by this specification.

#### 13.1.4.2 Fault LED Out

The Fault LED Out signal is driven by the drive when the drive detects an abnormal condition. The implementation of the Fault LED Out signal is optional.

The Fault LED Out signal is designed to pull down the cathode of a LED. The anode is attached to the proper +5 voltage supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive. Reference table 2 for the LED driver characteristics.

#### 13.1.4.3 -Write Protect

The Write Protect line signals the drive to disable write operations to the media. Drives are required to implement this feature. When the -Write Protect line is GND, writes are disabled. When the -Write Protect line is OPEN, writes are permitted. Reference table 3 for the definition of signal level.

#### 13.1.4.4 Reserved Signals

Reserved Signals shall have no electronic connection to the disk drive or motherboard circuitry until use for those signals is defined by this profile.

#### 13.1.4.5 Spindle Sync

The spindle sync is assigned a single pin, Spindle SYNC. The synchronization protocol and the electronic requirements for the SYNC signal are defined in the drive specification. Industry standards presently require that the drives interconnected for synchronization be the same or equivalent models. Spindle synchronization is managed by use of the SCSI Rigid Disk Geometry Mode Page (04h). The signal current requirements shall not exceed 24 milliamperes and the signal voltage shall not be higher than 5.25 or lower than -0.25 volts.

The SYNC signal is a source for noise and may be affected by noise. The design of the SYNC signal interconnections should take this into account by properly laying out the SYNC signals on the backplane or motherboard. Proper layout must consider routing relative to other signals, the proper line impedance, and terminations if necessary. The selection of the electronic transceiver must also take into account the possibility of noise. The signal levels, signal risetime, receiver thresholds, and receiver hysteresis must be considered as part of that selection.

#### 13.1.4.6 Active LED Out

The ACTIVE LED OUT signal is driven by the drive when the drive is performing a SCSI operation. The ACTIVE LED OUT signal is required to be implemented and is used to indicate that the disk drive is operating. Other optional indications can be provided by properly flashing the LED. The host system is not required to generate any visual output when the ACTIVE LED OUT signal is raised, but if such a visual output is provided, it must be white or green to indicate that normal activity is being performed.

The ACTIVE LED OUT signal is designed to pull down the cathode of an LED. The anode is attached to the proper +5 voltage supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive.

**Table 2. Output Characteristics of LED Driver Signal**

State	Current Drive Available	Output Voltage
Drive LED OFF	$0 < I_{OH} < 100 \mu\text{A}$	
Drive LED ON	$I_{OL} < -30 \text{ mA}$	$0 < V_{OL} < 0.8\text{V}$

#### 13.1.4.7 PBC Interlock CH1 - CH2

These signals control Port Bypass Circuits (PBC) located external to the disk drive. The PBC allows a loop to remain functional in the event of a drive failure or removal.

The PBC Interlock signal is generated under certain loop failure conditions within the disk drive, on the detection of an LPB primitive or on the removal of the disk. This signal causes an PBC circuit to actively bypass the drive. In the bypass state, the disk drive continues to receive on the inbound fibre. Another device on the loop may issue an LPE primitive addressed to the bypassed disk to bring it back into the loop. PBC Interlock CH1 controls the PBC in the loop connected to Channel 1. PBC Interlock CH2 controls the PBC on the alternate channel.

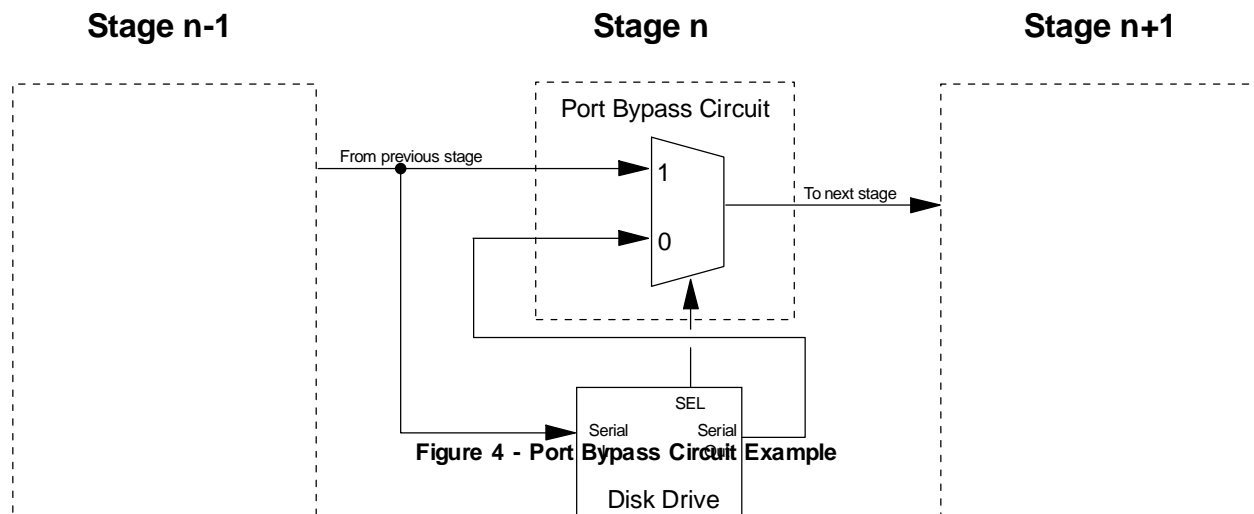
Failure modes that will enable the bypass circuitry as detected within the disk drive include:

1. Transmitter/ receiver wrap failure,
2. Loss of word sync,
3. Loss of read clock detect,
4. Loss of transmission clock detected,
5. Cyclic Redundancy Check (CRC) hardware failure,
6. 8b/10b hardware error, and
7. Protocol chip error.

An example of a Port Bypass configuration is shown in Figure 4 below.

After an LPB is received, the device shall:

1. go to MONITORING state until enabled by LPE or REQ(enable L\_Port)
2. not transmit on the Loop or attempt arbitration
3. set the P\_BYPASS state variable which activates the Loop Bypass Circuit, if present
4. does not propagate OPNyy or OPNyx addressed to it, but does propagate OPNfr, OPNyr, and OPNs to other ports
5. lose it's AL\_PA (goes to non-participating mode) if any form of LIP is received, but shall not go to OPEN-INIT state and shall not clear the P\_BYPASS state variable.



Note: Motorola MC10SX1189 is an example of a single-chip PBC.

### 13.1.4.8 Motor Start Controls

The method of starting the drive's motor is established by the signals RMT\_START and DLYD\_START. The state of these signals can either be wired into the backplane socket or driven by logic on the backplane. See table 4 for the definition of OPEN and GND.

**Table 3. Definition of Motor Start Controls**

Case	DLYD_START	RMT_START	Motor Spin Function
1	OPEN	OPEN	Motor spins up at DC power on.
2	OPEN	GND	Motor spins up only when SCSI "start" command is received.
3	GND	OPEN	Motor spins up after a delay of 12* seconds times the numeric SEL ID of the drive from DC power on.
4	GND	GND	Reserved. Drives not implementing this option shall execute power control according to the rules of Case 2.

\* This value may be reduced by drive suppliers to reflect the worst case time duration of peak current drains at the 12 volt or 5 volt source (or both) during motor spin up. In no case should the delay exceed 12 seconds.

**Table 4. Electronic Requirements for Inputs**

State	Voltage	Current
OPEN	$2.4 < V_{IH} < V_{CC} + 0.5$	$0 < I_{IH} < \pm 100 \mu A$
GND	$-0.5V < V_{IL} < 0.4V$	$0 < I_{OH} < -3 \text{ mA}$

### 13.1.4.9 Fibre Channel Signals

The Fibre Channel signals implement a standard Fibre Channel Arbitrated Loop protocol. Standard FC-0 voltage and current levels are supplied to the drives and expected from the drives. The signals are transmitted on transmission lines (Twinax) and capable of long distance (30 meters) and high data rates (up to 1.0625 GHz).

### 13.1.4.10 SEL\_6->SEL\_0 ID Lines

See FC-AL, Annex K for a description of how the 7 SEL lines are mapped to AL\_PAs( based on ascending arbitration priority).

The SEL\_6->SEL\_0 ID lines provide a binary decode associated with a physical unit for the full addressable logic range and beyond of the FC-AL. SEL decodes of 127 and 128 are recognized, but exceed the logical addressing range of the FC-AL. The SEL lines are used in a system to assign a physical location to a logical drive which greatly aids maintenance of a large array (finding the drive!) or to establish configuration control. The SEL lines may be thought of as the equivalent of a "backpanel logic plug". The SEL lines have absolutely nothing to do with the logical addressing of the drives unless the system so defines them that way. The system must handle the potential of duplicate SEL decodes in a large array.

If a drive were removed from the backpanel and plugged into a different location in the backpanel, the system normally uses the label information recorded on the drive to recognize where this drive should be mapped into its file management tables. This flexibility can be very desirable in data security applications or for configuration



expansion. In some RAID configurations, the physical location in a string of drives can be extremely important and should remain consistent. It is recommended that physical location dependent systems use the capability to read the SEL ID value and the logical address to insure proper configuration in the event of a maintenance or other error.

### 13.1.5 Backpanel Connector

The drive connector will mate with the AMP Champ 557189-1 vertical receptacle. The connector has been modified to allow for 0.040" of alignment tolerance along the X and Z axis. The connector will not be required to support the weight of the drive, although some vibration modes may place forces on the connector. Drives will be installed with appropriate brackets that will locate the drive and lock it in place.

### 13.1.6 Drive and Backpanel Connector Properties

The "ribbon" or "leaf" contact connector, AMP Champ 0.050" Series 1, meets the following requirements:

Pin count	40
Contact resistance	35 mΩ maximum
Contact current rating	1 A per isolated contact
Durability	1200 mating cycles minimum at maximum allowed misalignment 0.040" max
Insertion force	90 g/contact
Withdraw force	20 g/contact
Insulation resistance	1000 MΩ minimum at 250V DC; 500V DC desired
Dielectric withstanding volt	500V AC (rms), one minute
Hot plugging	Under Study
Contact plating	Appropriate to application requirements
Housing material	Appropriate to application requirements
Board retention	Provided by backplane and associated hardware

## 13.2 Direct Attach Cable Specifications

The internal twinaxial cable has two parallel wires surrounded by an aluminized polyester foil shield (and optionally a served shield). There may also be a drain wire within the shield (optional). The conductors in this cable may either be solid or stranded. Both constructions are good, but the stranded cable may be stronger for applications requiring more flexibility or physical abuse. Differential signal propagation skew must also be controlled if the application requires longer cable lengths (>15 meters). If the skew is controlled and the cable meets the loss specification (below) and the internal twinax electrical model, transmission of 1.0625 Gbaud Fibre Channel data at 30 meters is feasible.

**Table 5. Direct Attach Twinax Electrical Performance (Worst-Case)**

Frequency in Mhz	Attenuation (max.) per 30 meters	Velocity %C (max.)	Skew Differential
106	13.0 dB	78.99	< 200psec/30m
212	14.5	79.18	< 200psec/30m
318	16.0	79.32	< 200psec/30m
425	17.5	79.39	< 200psec/30m
531	18.6	79.43	< 200psec/30m
1590	26.5	80.50	< 200psec/30m

**Table 6. Direct Attach Twinax Electrical Characteristics (Worst-Case)**

Value	Tolerance	Units	Description		
+/- 0.5		pF/foot	Capacitance (sig-sig):	6.4	
		pF/foot	Capacitance (sig-gnd):	17.4	+/- 0.5
+/- 0.07		uH/foot	Inductance (sig-sig):	0.3	
+/- 0.04		uH/foot	Inductance (sig-gnd):	0.15	
	+/-5		Resistance (sig):		45
		mOhm/foot			
	+/-5		Resistance (gnd):		35
		mOhm/foot			

### 13.3 Direct Attach Cable Connectors

Industry standard 0.025" square posts on 0.100" centers form the basic receptacle function for the internal twinaxial connectors. These styles of connectors are available from numerous manufacturers and both polarized and non-polarized versions are available as the application requires. In addition, some forms of connector/cable retention have also been developed by the manufacturers.

Both shielded and un-shielded connector systems are recommended for internal FC applications. Shielded connectors are available from Gore as pre-terminated twinax assemblies (Gore P/N HSN3183-L, Rev. A) using their Shielded Interconnect Connector (SIC). Gore also has a cable assembly that features polarization and a positive locking mechanism (Gore P/N HSN3444-L). AMP has an unshielded connector/twinaxial cable system that offers polarization and a friction-lock method of cable retention (AMP P/N 93-4810-075-2).



## 14 Direct Attach Electrical Requirements

### 14.1 Direct Attach Transmitter Specifications

The transmitter shall be compliant with the FC-PH document and meet the jitter (both random and deterministic), frequency tolerances, output voltage, and rise/fall times. The transmitter specification is intended to be technology independent but delineates the minimum requirements to ensure compatible operation with FC-PH. The outputs were chosen to be ECL/PECL compatible and since they are AC-Coupled signals, that choice does not limit the technology as CML, Bipolar, GaAs, CMOS, etc. are all capable of the specified output levels. The output characteristics of the transmitter are defined in an eye diagram fashion in Figure 5.

**Table 7. FC-AL 1.0625Ghz Transmitter Specifications**

Frequency Tolerance ppm	+/- 100
Output Voltage Max/Min mV(p-p)	1600/600
Deterministic Jitter % (p-p)	10
Random Jitter % (p-p)	12
Rise/Fall Time (20-80%) max, ps	400



## 14.2 Direct Attach Receiver Specifications

The receiver shall be compliant with the FC-PH document and capable of recovering 8B/10B encoded data which complies with the timing tolerances and amplitude levels given in that document.

The receiver must be capable of acquiring "lock" within 2500 bit times for either a start-up condition or a "phase-jump" which can include even an abrupt 180 degrees phase jump. "Lock" is defined as the PLO or timing mechanism being within 5% of the final frequency and jitter tolerances of an infinitely stabilized receiver. Over-sampled styles of receivers are possible and may be applied to the FC-AL as well. Over-sampled receivers can achieve lock within one word (4 bytes) time.

**Table 8. FC-AL Receiver Specifications**

RECEIVER	1.0625 GHZ	531.25 MHZ	265.625 MHZ
Minimum Amp. D21.5 Idle, mV	200	200	200
Maximum Input mV(p-p)	1600	1600	1600
S <sub>11</sub> Input Parameter (0.1->1.5) x the bit rate)	-17 dB	-17 dB	-17 dB
Minimum Connector Return Loss (.3-1 Ghz)	15 dB	15 dB	15 dB

### 14.3 Equalizer Specifications

The equalizer is a circuit which corrects for the natural attenuation of higher frequency components of a signal during transmission down a frequency-limiting cable. Since the quality of the cable determines the amount of equalization needed, no one equalizer is necessarily optimum for all applications. The design goals of the Electrical Fibre Channel efforts have been to define an interface that has sufficient margin to run without an equalizer, but adding an equalizer adds to the error rate margins and the margins needed to withstand extraneous, real-world events like lightning strikes, earth currents, ESD, power, and temperature variations.

The passive equalizer is designed with two intents:

1. Correct the higher frequency signal component attenuation by reducing the gain of the lower frequency signal components.
2. Correct for signal dispersion by adjusting phase delay of signal components to correspond to the correct time-of-arrival as they were transmitted.

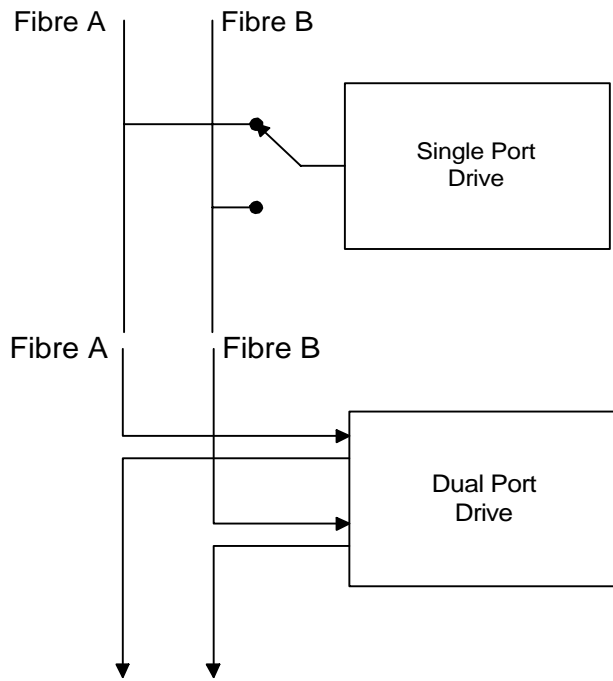
Practice has shown that optimizing the equalizer for about 2/3rds the length of the maximum cable length yields satisfactory results in terms of overall error rates. Such an equalizer then can also be used for twinax cables ranging from very short to long (30 meters) without compromising transmission error rates. Lower frequency FC-AL implementations may even go considerably beyond 30 meters as a maximum cable distance. Muratta-Erie has developed an equalizer (RLC-10015) in a small hybrid ceramic package which is optimized for a 75 ohm impedance and thus could be applied to a single side of the twinax signal. Two such equalizers would be required obviously for the twinax. Shortly, a single equalizer package should be available.

## 15 Dual Ports

Dual ports may be implemented to provide a redundant path for reliability or performance. The definition of dual port is outside the scope of this specification. Some examples of implementations are provided to illustrate the concept.

In this example, a single ported device accesses 2 fibre links by use of a serial switch. This switch could be on the drive or backplane.

In this example, a dual port drive is used. Both fibre links have receivers, transmitters, and some degree of protocol logic. It is the degree of protocol logic that determines how much functionality the ports have. One implementation may let the first port selected have access to the FC2 resources while the other port returns a CLOSE primitive to an OPEN or just issues no credit. Another implementation may allow the second port to receive commands while the first active port is doing a data transfer.



**Figure 7 - Dual Port Implementation Example**

## 16 External Electrical and Mechanical Requirements

The Small Form Factor committee which meets during X3T10 weeks will adopt standardization of external Connector and Transceiver characteristics for FC-AL bulkheads. Contact: Dal Allan, 916-867-6630, 2501752@mcimail.com.

## Annex A: Pending ANSI Standards Changes

### A.1 FC-PH/EP

#### A.1.1 Class 3 Streamed and Open Sequence Definitions

There are several places in FC-PH where the terms "Streamed Sequence" or "Open Sequence" are used or defined inconsistently or inappropriately, particularly for class 3:

Glossary (3.1.166)	SEQ_ID (18.6)	Active and Open Sequence (24.6.1)
<p><i>"streamed Sequence: ...Any new Class 3 Sequence initiated before the expiration of R_A_TOV for all Data frames in the previous Sequence (see 18.6)."</i></p>	<p><i>"If the Sequence Initiator initiates a new Sequence for the same Exchange...before R_A_TOV has expired for all frames of a Class 3 Sequence, it is termed a streamed Sequence."</i></p>	<p><i>"In Class 3, the Sequence Initiator considers the Sequence Open until the deliverability is confirmed or an R_A_TOV timeout period has expired."</i></p>

Reusing a SEQ\_ID after R\_A\_TOV expiration at the Sequence Initiator is a sufficient but not a necessary condition for ensuring that Sequences are not streamed in class 3. For example, ULP interlocks such as in FCP (XFR\_RDY, FCP\_RSP) which pass Sequence Initiative within an Exchange acknowledge completion of Sequences and allow SEQ\_ID reuse without waiting R\_A\_TOV.

Clauses 3.1.166 and 18.6 attempt to make class 3 "reliable" at the link layer through the use of inefficient timeouts, when ULP interlocks are a more appropriate method of achieving this end. Clause 24.6.1 hints at the use of these interlocks, and in the process contradicts 3.1.166 and 18.6. In order to unify the concepts of open Sequences and streamed Sequences for **ALL** classes of service, I suggest the following changes:

- a) Replace definition of "streamed Sequence" in 3.1.166 with:  
*"In any class of service, when a Sequence Initiator initiates a new Sequence for an Exchange while it already has open Sequences for that Exchange, the new Sequence is termed a streamed Sequence (see 24.6.1)."*
- b) Replace sentence 1, ¶4 in 24.1 with:  
*"FC-PH allows an N\_Port to initiate a new Sequence for the same Exchange while it already has Open Sequences for that Exchange."*
- c) SEQ\_ID (18.6), replace sentence 1, ¶2 with:  
*"In any class of service, if the Sequence Initiator initiates a new Sequence for an Exchange while it already has open Sequences for that Exchange, the new Sequence is termed a streamed Sequence (see 24.6.1)."*
- d) Add to the end of 24.6.1, ¶1:  
*"The determination of deliverability of class 3 Sequences is beyond the scope of FC-PH, which provides no deliverability guarantees for class 3 Sequences."*
- e) Delete the following from 24.3.8f (It is not true, and deliverability of class 3 Sequences is beyond the scope of FC-PH):  
*"In Class 3 a Sequence Initiator shall consider a Sequence as deliverable and complete only if it has received a Basic Accept in reply to an ABTS frame or if it has received an Accept to a Read Exchange Status request which confirms the deliverability."*

## **A.1.2 Fabric Login Required (23.2, p171)**

Clause 23.2 states,

*"Login with the Fabric is required for all N\_Ports"*

The intent was to allow ports which support both point-point and fabric topologies to have a single FLOGI mechanism for discovering their topology (the ACC indicates N\_Port or F\_Port).

However, this is unnecessarily restrictive for devices which can only exist in certain topologies or which do not have microcode to perform FLOGI. For example, Target-only NL\_Ports on private loops do not have the capability nor the need to perform FLOGI. Also, early implementations which can only communicate point-point and cannot yet perform FLOGI should still be considered compliant (just as early fabrics cannot communicate F\_Port to F\_Port but are still considered compliant).

Therefore, this clause should either be stricken, or modified to read:

*"Login with the Fabric is required for all N\_Ports which support Fabric topologies."*

## **A.2 FCP**

### **A.2.1 Task Management and FCP\_RSP**

New PRLI parameter which specifies whether or not:

- Task Management is sent using I1 instead of T5, and
- FCP\_RSP (I4) is transmitted in response to I1 Task Management
- which codes to use if Task Management successful, unsuccessful, or unsupported

## **A.3 SAM (Rev 16)**

*[ed note: the below is still under discussion]*

### **A.3.1 ACA ACTIVE FAULTED INITIATOR status**

In SAM section 6.2, add the following after ACA ACTIVE.

ACA ACTIVE FAULTED INITIATOR. This status shall be returned to the faulted initiator when an auto contingent allegiance exists within a task set for the faulted initiator and the faulted initiator issues a command for that task set when at least one of the following is true:

- a) There is a task from the faulted initiator with the ACA attribute in the task set, or
- b) The task created to execute the command did not have the ACA attribute and the NACA bit was set to one in the CDB control byte of the faulting command (see clause 6.6.1).

In Section 6.6.1.1, paragraph 3, append the following sentence:

While an ACA condition is in effect, tasks created by the faulted initiator which are closed with a status of ACA ACTIVE FAULTED INITIATOR shall not be entered into the faulted task set.

In Section 8, paragraph 3, sentence 2, add the new status to the list of statuses which prevent a task from being entered into the task set:

A task shall be entered into the task set unless a condition exists which causes that task to be completed with a status of BUSY, RESERVATION CONFLICT, TASK SET FULL, ACA ACTIVE, ACA ACTIVE FAULTED INITIATOR, or CHECK CONDITION...



Section 6.6.4, first sentence second paragraph. In order to clarify what is meant by "transferred", change to:

Sense data shall be preserved by the logical unit for the initiator until the initiator has confirmed receipt of the sense data or until another task from that initiator is entered into the task set.

## Annex B: Nonzero Login\_BB\_Credit

This Annex describes how to interoperate using nonzero Login\_BB\_Credit on FC-AL. While early implementations which conform to the Private Loop Profile may only interoperate using Login\_BB\_Credit = 0, enhanced versions may allow the OPN initiator to transmit frames without waiting for R\_RDYs from the OPN Recipient.

### B.1 Terminology

**Login\_BB\_Credit:** On FC-AL, equal to the number of receive buffers that a receiving L\_port shall guarantee to have available when a circuit is established. Login\_BB\_Credit is discovered in the login protocol.

**Available BB\_Credit:** A variable used by a transmitter to determine permission to transmit frames, and if so how many. It is incremented upon receipt of R\_RDY, decremented upon transmission of a frame, and gives the transmitter permission to send another frame as long as it is positive. On transmission of OPN, there are two methods which the transmitter can receive Available BB\_Credit:

- 1) implicitly on every OPN via the Login\_BB\_Credit advertised in login by the OPN Recipient, and
- 2) explicitly on every OPN via R\_RDYs received from the OPN Recipient

### B.2 OPN, CLS, and BB\_Flow Control Rules

- 1) When Login\_BB\_Credit > 0 at the OPN Initiator, it shall open **half duplex**. This is to prevent buffer overrun if a CLS is originated at the same time that data frames are in flight on the return path from the CLS recipient, followed by an immediate open from another NL\_Port with up to Login\_BB\_Credit more frames arriving before the buffers are available.

*[ed note: for forward compatibility, first generation devices must be able to be opened half duplex. Options are:*

- 1) *no change, no compatibility - Login\_BB\_Credit > 0 not interoperable,*
- 2) *change profile so that first generation devices may be opened either full or half*
- 3) *change profile so that all devices may be opened half-duplex in all generations of private loops]*

- 2) When Login\_BB\_Credit > 0 at the OPN Recipient, it shall transmit no fewer than Login\_BB\_Credit R\_RDYs back to the OPN Initiator when opened. This provides compatibility with OPN Initiators which wait for R\_RDYs before transmitting.
- 3) When Login\_BB\_Credit > 0 at an NL\_Port, that NL\_Port shall neither originate nor forward a CLS until it can guarantee that the number of available receive buffers on an immediate, subsequent OPN will be greater than or equal to its Login\_BB\_Credit.
- 4) If an OPN Initiator chooses to take advantage of nonzero Login\_BB\_Credit advertised by an OPN Recipient, it must be aware of the value of Login\_BB\_Credit at the OPN Recipient without receiving any R\_RDYs from it. The OPN Initiator shall transmit up to Login\_BB\_Credit frames immediately following the OPN (half duplex) without waiting for R\_RDYs, and **discard** the first Login\_BB\_Credit R\_RDYs that arrive from the OPN Recipient, regardless of how many frames were transmitted (in order to avoid "double crediting" the OPN Recipient which must observe rule 1).

## **Annex C : Changes from Previous Revisions**

### **C.1 Version 1.20 to Version 1.30**

- X\_ID interlock removed (only pertains to class 1 and 2 - see FC-PH definition)
- Frame size = max on all frames but the last frame of a Sequence
- PRLI: only those pages returned in ACC which are present in the PRLI request
- Reset table: note 9 added
- CMD, XFR\_RDY, and RSP required to be single-frame Sequences
- FCP\_DL requirement elaborated
- Task management responses elaborated
- SEQ\_ID reuse clarified
- FC-PH/EP proposal in Annex A updated
- Compatibility issue with Login\_BB\_Credit highlighted in Annex B
- Annex C added