



BROCADE

May 20, 2001

T10/01-197 Revision 0

To: John Lohmeyer, Chair, T10 Standards Committee, NCITS
From: Robert Snively
Date: June 20, 2001
Subject: Resolution of FCP-2 public review comment

1 Introduction

A public review comment was received from Roger Cummings of Veritas Software. The public review comment addressed possible alternative requirements for clearing current MODE SELECT pages and reservations for tape drives. The comment is contained in document T10/01-176r0.

2 Proposed resolution

The public review comment is rejected for the reasons described below.

3 Summary of the public review comment

The public review comment proposes two changes to FCP-2 table 4, "Clearing effects of link related functions".

3.1 Requested change to MODE SELECT / MODE SENSE behaviors.

The first change is intended to maintain current mode page settings of tape drives over link failure conditions that may be recoverable by some drivers.

The change proposes separating the consideration of LOGO/PLOGI into two classes of behavior:

- Explicit Logout

- Implicit Logout and PLOGI

The text of the proposal is ambiguous about PLOGI. I interpreted the second item to include all cases of PLOGI, both implicit and explicit, but to include only implicit logout for other logouts. Either interpretation would not change the resolution of the review comment.

All rows remain unchanged for both cases except for the row concerning MODE SELECT/MODE SENSE page behaviors. For this row, the row is subdivided into two rows, as follows:

- Only for initiator port associated with the action (random)

- Only for initiator port associated with the action (serial)

It is assumed in this analysis that the words "serial" and "random" were probably meant to apply to devices compliant with the SCSI Stream Commands (SSC-2) document and the SCSI Block Commands (SBC-2) document respectively. Note that either type of device may have optional serial or random behavior that is not known at the SCSI level.

The change makes no modification to the behavior of block command devices. The change requires stream command devices (typically tape drives) to retain the current mode pages

instead of modifying their behavior to the saved or default mode pages for the following conditions:

- Failed Discovery after LIP
- Failed Discovery after OLS
- Implicit Logout & PLOGI

3.2 Requested description of stream command model in FCP-2

The second change is intended to prevent link recovery operations from modifying the state of tape drives.

A specific note is requested to indicate that stream command devices shall not lose or alter media positioning information or perform a rewind for any link related function except an explicit logout (LOGO). The note would be included in the footnotes of Table 4.

3.3 Requested change to behavior of reservation operations

The third change is intended to maintain legacy reservations of tape drives unmodified by any link operation except explicit LOGO.

The change proposes separating the consideration of LOGO/PLOGI into two classes of behavior as described in 3.1:

- Explicit Logout
- Implicit Logout and PLOGI

All rows remain unchanged for both cases except for the row concerning Device Reservation behaviors. For this row, the row is subdivided into two rows, as follows:

- Only for initiator port associated with the action (random)
- Only for initiator port associated with the action (serial)

It is again assumed in this analysis that the words “serial” and “random” were probably meant to apply to devices compliant with the SCSI Stream Commands (SSC-2) document and the SCSI Block Commands (SBC-2) document respectively.

The change makes no modification to the behavior of block command devices. The change requires stream command devices (typically tape drives) to retain their reservations unchanged under the following conditions:

- Failed Discovery after LIP
- Failed Discovery after OLS
- Implicit Logout & PLOGI

4 Review of relevant standards related to comment

4.1 Review of proposed change to MODE SELECT / MODE SENSE behaviors

The following text refers to the behavior of MODE SELECT/MODE SENSE pages under various conditions.

FC-PLDA text:

Subclause 10.5, table 16

The text indicates that LOGO and PLOGI shall restore mode pages for the port initiating the action to their saved or default states. No distinction is made for tapes, even though tapes are considered in other portions of the document.

FC-TAPE text:

Subclause 7.2.2

The text describes the allowed properties of the Enable Modify Device Pointers (EMDP) bit and the burst length field. The host adapter restoring a link after a reset must assume a set of properties for both fields during the restoration of the link. This can only be achieved if those parameters are restored to a saved or default value. Failure to agree upon these values may cause protocol errors and data integrity errors.

SPC-2 Text:

Subclause 7.11.2

A PC field value of 00b requests that the device server return the current values of the mode parameters. The current values returned are:

- a) the current values of the mode parameters established by the last successful MODE SELECT command;
- b) the saved values of the mode parameters if a MODE SELECT command has not successfully completed since the last power-on or hard reset condition; or
- c) the default values of the mode parameters, if saved values, are not available or not supported.

This clearly defines the correct value of the mode parameters after a clearing event like the loss of a login and therefore a process login image pair. This is not new text, but dates back to SCSI-2 (see 8.2.10 of SCSI-2). Any device not following this and any program not expecting this behavior is operating in a non-standard manner.

4.2 Requested description of stream command model in FCP-2

The following text refers to the clearing of command state during various link operations.

FC-TAPE text:

Clause 5.2, 5th paragraph:

The text indicates the following:

SCSI Initiators and SCSI Targets shall validate the current Fabric login following every Loop Initialization by comparing the Loop Fabric Address, Fabric Port Name, and Fabric Name received during FLOGI with those reported by the received FAN. All three identifiers reported by the FAN shall match the values reported during FLOGI or a configuration change has occurred; LOGO is required and all open Exchanges shall be terminated.

This termination of all exchanges leaves the state of the device uncertain, since there is no way to determine which commands have been completed and which have not.

SSC-2 text:

Many clauses:

Many clauses indicate that various machine states shall be reset or set to a default value by a reset or power off sequence. As a minimum, Unit Attention shall be presented. This would be the expected behavior in the case of any potential reconfiguration, which includes all cases of implicit LOGO and mismatched FC-PLDA or FC-FLA authentication attempts.

4.3 Requested change to behavior of reservation operations

The following text refers to the clearing of reservations during various link operations.

FC-TAPE text:

Clause 9.5.1, Table 25

All FC-TAPE compliant devices are required to implement persistent reservation. Persistent reservation was designed to be maintained across all resets and even power off conditions. The Reserve/Release requirements do not include third-party reservations, required for LAN-free backup procedures.

SPC-2 text:

Clause 5.5.2, paragraph 2

Reservations managed using the Reserve/Release method do not persist across some recovery actions (e.g., hard resets), so most systems require significant re initialization after a failure that results in a hard reset. Reserve/Release managed reservations are retained by the device server until released or until reset by mechanisms specified in this standard.

This text clearly indicates that those failures that result in a hard reset will cause legacy reservations to be reset. There is no reset harder than the loss of a login relationship, which automatically removes the process login image pair between an initiator and a target.

Clause 5.6

This clause describes the actions that clear reservations, including TARGET RESET, LOGICAL UNIT RESET, and CLEAR TASK SET. Power off and hard reset will also clear reservations.

5 Resolution of comment

5.1 Requested change to MODE SELECT / MODE SENSE behaviors.

By reviewing the FC-PLDA, FC-TAPE, and SPC-2, it is clear that hard resets are intended to return the current mode page settings of a SCSI device to the saved or default values. FCP and FCP-2 have intended to maintain that characteristic. FCP-2 is the first SCSI link standard to use the rigorous tabling mechanism pioneered by FC-PLDA to clarify that behavior, but it has made no change to the already standard behavior. Those implementations that do otherwise are working outside the SCSI standards family.

The saved values are established by the initiator such that the correct values will be present after a hard reset. An initiator that cannot save values and cannot operate with the default values should update the mode page values to the desired values during the initialization process. The requirement for the initialization process will be indicated by the requirement for a new login, process login, and the presentation of UNIT ATTENTION conditions.

Note that partition information and certain other MODE SENSE information defined by SSC-2 may remain unmodified by hard resets unless the medium is removed.

Changing FCP-2 as requested by the ballot comment would cause it to violate FC-PLDA, FC-TAPE, SPC-2, and would obsolete a long line of compliant implementations. For this reason, this part of the ballot comment is not accepted.

5.2 Requested description of stream command model in FCP-2

By reviewing FC-TAPE and SSC-2, it is clear that device state is not required to be maintained across reset actions. There is no reset action more severe than a reconfiguration which has not been verified to contain the same pair of devices. Such a reconfiguration requires complete analysis of the present state of all devices and the execution of appropriate recovery algorithms. Fortunately, all FC-TAPE compliant devices are required to implement the LOCATE commands, facilitating recovery.

Note that failures that are temporary and do not indicate a change in configuration are not resetting actions and are recovered normally.

Changing FCP-2 to mandate certain behavior specific to tape drives would place new requirements on FC-TAPE and SSC-2 compliant devices and make such devices obsolete. For this reason, this part of the ballot comment is not accepted.

5.3 Requested change to behavior of reservation operations

Persistent reservations were explicitly designed to operate successfully across both reset actions and power off actions. The RESERVE and RELEASE commands were never upgraded to allow such behavior. All FC-TAPE compliant devices are required to implement persistent reservations. The problem described in the ballot comment is best resolved by using the command set invented specifically for that purpose.

Changing FCP-2 to modify the behavior of RESERVE and RELEASE would obsolete present implementations that are compliant with SPC-2. For this reason, this part of the ballot comment is not accepted.