

May 20, 2001

To: Deborah Donovan, NCITS Secretariat
Cc: PSA Dept., ANSI
John Lohmeyer, T10 Chairman
Robert Snively, FCP-2 Technical Editor
From: Roger Cummings, VERITAS Software
Subject: Public Review Comment on dpANS NCITS 350, FCP-2

Please accept this document as a Public Review comment by VERITAS Software against NCITS 350, Information Technology - Fibre Channel Protocol for SCSI, Second Version (FCP-2).

VERITAS will be pleased to have representation at the next meeting of Technical Committee T10 to discuss this comment, and the changes suggested herein, in detail.

Regards,

Roger Cummings
VERITAS Software
roger.cummings@veritas.com
407.531.7257

PUBLIC REVIEW COMMENT on NCITS 350 FCP-2

Introduction

VERITAS has decided to take the serious step of submitting a Public Review comment on this draft standard because of experience gained in testing prototype equipment that conforms to the requirements of the FCP-2 draft along with our data protection applications. This experience has led us to believe that there is a serious problem with two of the definitions contained in FCP-2. We believe that these problems will be of significance to many vendors such as VERTAS who produce backup applications that make use of serial access storage devices. If left uncorrected, the problems will require those applications to deal with FCP-2-compliant devices quite differently than with existing devices that conform to the previous version of the SCSI-3 Fibre Channel Protocol (FCP) definition (X3.269:1996), or to any of the SCSI-3 Parallel Interface (SPI) standards and their predecessors. Specifically, we have experience with devices in the field that are proven compatible with our software and compliant with these standards but which cease to be compatible when new firmware which meets the requirements of FCP-2 is installed.

Background

Backup applications are mission-critical to many enterprises. The applications run for extended periods of time, often many hours, and because of their impact on system performance and data availability they are often constrained to run during specific time “windows” when system loading is less than critical and when, where required, operators are available to supervise and handle media. In many cases the backup application runs for the entire duration of each window.

If a system problem, such as a hardware failure or a configuration change, is encountered during the execution of the backup it is therefore impractical to restart the entire process because it would not be able to complete during the window. Because of this, most backup applications incorporate sophisticated logic, developed over many years, to be able to salvage the uncompleted backup and complete it without requiring a complete restart. Schemes such as the disconnection of the serial access drive and reconnection to a different channel or switch port are supported by this logic. The ability to perform such a recovery of a half-completed backup is regarded as a key capability by many users of backup applications. It is therefore part of the testing that is performed before the release of any new revision of a backup application.

Such a recovery scheme depends on the ability of the serial access drive to maintain a significant amount of “state” information across the physical disconnection and reconnection. Much of this state information is set and monitored through SCSI mode pages. Specific examples of such information include the media position, compression or encryption scheme used, and even the block size used for fixed-block media formats. All existing SCSI serial access devices, including those that conform to X3.269, FCP, support the maintenance of such state information. In the case of FCP the effect of link activities on SCSI mode pages and other state information is undefined.

Problem

Table 4 on page 16 of the NCITS 350, FCP-2 Public Review draft defines the clearing effect of link related functions. The seventh section of this table defines that:

Target mode page parameters restored from saved pages (when saved pages are supported, or default mode page parameters when saved pages are not supported)

for the initiator port associated with the action, for the following actions:

Failed Discovery after LIP
Failed Discovery after OLS
LOGO & PLOGI

Note 13 to Table 4 specifically states that LOGO may either be implicit or explicit.

By the definitions contained in X3.230:1994, the Fibre Channel Physical and Signaling Interface, and its successors, loss of signal (light) at the Fibre Channel interface, or the receipt of PLOGI is regarded as an Implicit Logout.

Therefore the definition quoted from Table 4 in concert with the definitions contained in FC-PH will cause the state information contained in a serial access device to be destroyed when it is disconnected, and will prevent the backup recovery strategy described above from succeeding. It will also cause that state information to be destroyed in response to a number of types of Fibre Channel reconfiguration and to an error recovery process (that reissues PLOGI) in the SCSI Initiator associated with the backup application.

Similarly the ninth section of Table 4 on page 16 defines that:

Device Reservations

are cleared for the Initiator Port associated with the action, for the same actions as identified above. This definition will also cause reservations placed on a serial access device to be cleared when it is disconnected, and will raise the issue of possible data corruption in the backup strategy described above.

Suggested Remedies

VERITAS understands that the contents of Table 4 of FCP-2 describe definitions that have been in place since the publication of NCITS TR-19:1998, Fibre Channel Private Loop Direct Attach (FC-PLDA), and have been carried forward in other Technical Reports. We believe that these definitions are appropriate for random block access devices, but not for serial access devices.

Therefore we request the following changes to FCP-2. The changes are intended to provide the same level of robustness in the use of serial devices that meet the FCP-2 definitions as is found with existing SCSI and FCP-compliant serial access devices. The requested changes are:

- 1) The current column in Table 4 which describes the response to LOGO and PLOGI be split into two columns – one for explicit LOGO only, and the other for Implicit Logout & PLOGI;
- 2) The explicit LOGO column has the same contents as the existing column;
- 3) The Implicit Logout & PLOGI column has the same contents as the previous column, except for the seventh and ninth sections (i.e. the section containing “target mode page parameters” and the section containing “device reservations”) of the table which is rewritten as follows;
 - a. The lines containing “Only for the initiator port associated with the action” are split into two lines, namely one for random access devices, and the other for serial access devices;
 - b. The lines for random access devices have the same contents as the existing line;

- c. The lines for serial access devices have a “-“ in the columns for Failed Discovery after LIP, Failed Discovery after OLS, and Implicit Logout & PLOGI;
- 4) Note 13 to Table 4 is deleted.
- 5) A specific requirement is included in subclause 4.10 (Port login/logout) or elsewhere to the effect that for serial access devices no link related function other than an Explicit Logout (LOGO) shall cause any target mode page parameters to be restored from saved pages (when saved pages are supported, or default mode page parameters when saved pages are not supported), device reservations to be cleared, media positioning to be lost or altered, or a rewind operation to commence.

A revised Table 4 in line with the changes above is shown below:

Table 4 - Clearing effects of link related functions

Target object	FC link action affecting FCP target object								
	Target Power Cycle	Reset LIP(y,x) ³	Failed Discovery after LIP ¹⁴	Failed Discovery after OLS ¹⁵	Implicit LOGO, PLOGI	Explicit LOGO ¹	PRLI ⁸ PRLO ¹³ , ¹	TPRLO ⁵	ABTS (Exchange)
PLOGI parameters	Y	Y	N	N	N	N	N	N	N
For all logged-in initiator ports	-	-	Y	-	Y	Y	N	N	N
Only for initiator port associated with the action	-	-	Y	-	Y	Y	N	N	N
Open FCP Sequences Terminated	Y	Y	N	N	N	N	N	Y	N
For all initiator ports with open FCP Sequences	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action	-	-	-	-	-	-	-	-	Y
Only for FCP Sequences associated with Aborted FCP Exchanges	-	-	-	-	-	-	-	-	Y
Login BB_Credit_CNT	Y	Y	-	-	-	-	N	N	N
For all Logged-In L_Ports	-	-	Y	Y	Y	Y	N	N	N
For transmitting L_Port only	-	-	Y	Y	Y	Y	N	N	N
Hard Address Acquisition Attempted	Y ¹	Y ¹	N	N	N	N	N	N	N
PRLI parameters cleared	Y	Y	N	N	N	N	N	N	N
For all logged-in initiator ports	-	-	Y	Y	Y	Y	Y	N	N
Only for N_Port or L_Port associated with the action	-	-	Y	Y	Y	Y	Y	N	N
Open Tasks (FCP Exchanges) Aborted	Y	Y	N	N	N	N	N	Y	N
All tasks for all initiator ports with open tasks	-	-	Y	Y	Y	Y	Y	-	N
All tasks, only for initiator port associated with the action	-	-	-	-	-	-	-	-	Y
Only for specified task	-	-	-	-	-	-	-	-	Y
Target mode page parameters restored from saved pages (when saved pages are supported, or default mode page parameters when saved pages are not supported)	Y	Y	N	N	N	N	- ¹²	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	- ¹²	-	N
Only for initiator port associated with the action (random)	-	-	-	-	-	Y	-	-	-
Only for initiator port associated with the action (serial)	-	-	-	-	-	Y	-	-	-
Pre-existing ACA, Unit Attention ⁷ , and Deferred error conditions cleared	Y	Y	N	N	N	N	N	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action	-	-	Y	Y	Y	Y	Y	-	N
Device reservations	Y	Y	N	N	N	N	N	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action (random)	-	-	-	-	-	Y	-	-	-
Only for initiator port associated with the action (serial)	-	-	-	-	-	Y	-	-	-
Persistent device reservations ¹⁰	Y ¹¹	N	N	N	N	N	N	N	N
For all initiator ports	-	-	N	N	N	N	N	-	N
Only for initiator port associated with the action	-	-	N	N	N	N	N	-	N
CRN (Command Reference Number) (set to one)	Y	Y	N	N	N	N	N	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action	-	-	Y	Y	Y	Y	Y	-	N
Prevent Allow Medium Removal state cleared to allow removal	Y	Y	N	N	N	N	N	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action	-	-	Y	Y	Y	Y	Y	-	N
Buffered data for XOR, EXTENDED COPY, COPY	Y	Y	N	N	N	N	N	Y	N
For all initiator ports	-	-	Y	Y	Y	Y	Y	-	N
Only for initiator port associated with the action	-	-	Y	Y	Y	Y	Y	-	N
Access controls data	N	N	N	N	N	N	N	N	N
AccessID enrollment state to pending enrolled	Y ¹⁶	Y	N	N	N	N	N	Y	Y
For all SCSI initiators in enrolled state	-	-	Y	Y	Y	Y	Y	-	-
Only for SCSI initiator port initiating action in enrolled state	-	-	Y	Y	Y	Y	Y	-	-