

July 5, 2001

To: John Lohmeyer, T10 Chairman  
Robert Snively, FCP-2 Technical Editor  
From: Roger Cummings, VERITAS Software  
Subject: Comments on the Response to the VERITAS Public Review Comment on  
dpANS NCITS 350, FCP-2

John,

VERITAS would like to thank Robert Snively for his consideration of our Public Review comment against FCP-2 (01-176r0), and his detailed response (01-197r0).

VERITAS has comments on the response, which are listed below. We are requesting an agenda item of 30 minutes in length at the next meeting of the SCSI Commands Architecture and Protocols Working Group to discuss the response and these comments in detail. Subsequent to the presentation, VERITAS will be pleased to accept the recommendation of the CAP WG, and the decision of the subsequent Plenary meeting of Technical Committee T10, on this subject.

Regards,

Roger Cummings  
VERITAS Software  
[roger.cummings@veritas.com](mailto:roger.cummings@veritas.com)  
407.531.7257

**VERITAS COMMENTS ON RESPONSES TO PUBLIC REVIEW COMMENT on NCITS 350 FCP-2 contained in 01-197r0.**

1) We apologize for the confusion as stated in sections 3.1 and 3.3 of the response regarding the words “random” and “serial”. We had hoped that the words in the first paragraph under “Suggested Remedies” in 01-176r0 would provide sufficient explanation, clearly they did not. However the interpretation of “random” and “serial” in the response is in line with our intentions.

2) Section 3.1 of the response comments on an ambiguity of the proposal with respect to PLOGI, namely that the behavior with respect to Implicit and Explicit PLOGI is not defined separately. This ambiguity is also present in FCP-2, where the clearing actions tables do not differentiate between the two forms of PLOGI. In fact implicit PLOGI is only mentioned once in FCP-2, in subclause 12.6, and there the behavior is identical for both forms. The authors of the Public Review comment were therefore working under the assumption that the behavior for both forms of PLOGI should be identical.

3) Section 4.1 of the response, under FC-TAPE text, states:

*“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.”*

No such requirement exists as stated in FC-TAPE, and we respectfully suggest that there is another way that host adapter and tape can assume compatible values for both fields, namely that they remain unchanged from the values previously established.

4) Section 4.1 of the response, under the text quoted from SPC-2, states:

*“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”.*

With respect, we believe it does not. The SPC-2 text (and the equivalent text in predecessor documents dating back to SCSI-2) describes the behavior after a hard reset and a power-on condition. The Public Review comment addresses neither situation.

5) Section 4.3 of the response, under FC-TAPE text, states:

*“All FC-TAPE compliant devices are required to implement persistent reservation.”*

However subclause 8.4 of FC-TAPE states:

*SCSI Targets that do not have access to non-volatile storage are not required to support Persistent Device Reservations across power cycles (i.e. APTPL=1).*

- 6) In similar vein to 5) above, section 4.3 of the response, under the text from quoted from SPC-2, states:

*“This text clearly indicates that those failures that result in a hard reset will cause legacy reservations to be reset.”*

However the text in subclause 5.5.2 of SPC-2 states:

*“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.”*

The references to “some recovery actions” and “most systems” should be noted. Again, however, this text is associated with hard resets only.

- 7) The first sentence of section 5.1 of the response states:

*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.*

In the first paragraph under Suggested Remedies, the Public Review comment acknowledged that the actions defined in Table 4 of FCP-2 have been in place since the publication of NCITS TR-19:1998, Fibre Channel Private Loop Direct Attach (FC-PLDA). (Note that FC-TAPE makes a forward reference to FCP-2 in this area.) However, we do not believe that this characteristic is defined in FCP. Subclause 6.2.5 of FCP defines that:

*Immediately after the execution of a PRLI, both members of the image pair shall have the same state as they would have after a hard reset or a power on with respect to each other.*

No such definition exists for PLOGI, We believe that this is why a number of the tape products in our lab do not treat PLOGI, or removal/insertion cycles of the physical interface, as hard resets to the SCSI command set “layer”.

Therefore FCP-2 represents the first time that the clearing actions of link related functions other than PRLI and PRLO will be specifically associated with SCSI command set “layer” actions in a T10 standard.

Note that the Public Review comment against FCP-2 does not seek to modify the behavior in response to either PRLI or PRLO, only PLOGI & Implicit Logout.

8) Section 5.3 of the response states:

*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.*

However such reconfigurations also have to be survived by persistent reservations, and methods of doing just this have been suggested in a number of T10 documents, including 98-206r1. Specifically, a comparison between the WWN in a new PLOGI allows a device with a previously-stored value allows a device to associate existing tasks with the device logging in, even if a physical reconfiguration has occurred. 98-206r1 states that:

*“FC-PLDA implies that ID and WWN relationships are remembered so that device verification can be performed in the event of an initialization”*

Surely this same mechanism could be used to associate device state across all FC link actions.

9) As far as VERITAS is aware, there is no other SCSI Protocol or Physical Layer standard or draft standard that mandates a change in the contents of mode pages in response to a physical layer action. As an example, Annex D of SPI-3, which describes the insertion and removal of SCSI devices in great detail, only mandates that all I/O processes for all SCSI devices be quiesced. This is in line with the definition of PLOGI (explicit or implicit) and logout (LOGO or signal loss) in the FC standards, which describe all outstanding Exchanges as being terminated.

## SUMMARY

The VERITAS position with regards to the Public Review comment and its response can be summarized as follows:

- a) FCP-2 is the first T10 dpANS to include detailed definitions of the clearing actions of link related functions, and therefore to link those functions with parts of the power on and hardware reset definitions in other SCSI standards. This is a major change from FCP.
- b) Insertion and Removal of devices is not directly addressed in FCP-2. The Fibre Channel standards define login and logout as a reset for the FC port, and the termination of all outstanding Exchanges. The latter part of this definition is consistent with the insertion and removal definitions in other SCSI standards. The only further definition of impact to the SCSI command set “layer” by physical level actions is contained in the T11 TR FC-PLDA.
- c) The vast majority of devices in the field that conform to the FC-PLDA are disks, and neither FCP-2 nor the Public Review comment seeks to change their defined behavior.
- d) We believe that T10 should decide if it is appropriate going forward that the operation of the SCSI Command Set “layer” should be so largely impacted by physical interface actions, with specific reference to the type of long-duration, time-critical task sets that are outlined in the Public Review comment. It should also be noted that some of these physical layer actions may be originated by third parties. If T10 decides that such a level of impact is no longer acceptable, then a change of the type proposed by VERITAS in the PR comment may be a sensible first step towards a new separation of functions.
- e) We believe that T10 should also be starting to consider the impact of transport mechanisms such as TCP/IP over a Wide Area Network on such definitions such as clearing actions at the SCSI Command Set layer. Ideally, from the view of an application vendor such as VERITAS, equivalent definitions should be created for all transports.