

To: T10 Membership
From: Lawrence J. Lamers, Adaptec, Inc. <ljlamers@ieee.org>
Subject: READ/WRITE BUFFER Behavior During Domain Validation
Date: Thursday, November 04, 1999

Background: The domain validation process needs some modifications to SPC-2 in order to function in an efficient manner.

In multi-initiator environments, there are issues with the use of RESERVE and RELEASE that cause problems for upper-level software getting a RESERVATION CONFLICT status, particularly in a multi-initiator environment. Linked commands are not supported by all targets and may not be supported in lower-level software. Furthermore linked commands have not been used much and the resulting side effects are indeterminate.

One alternate is to require separate echo buffers on a per initiator basis. This could waste resources as it is likely not implemented in all first generation Ultra160 devices, and may not be feasible on tape drives/CD-ROMs.

The solution proposed is for targets to report CHECK CONDITION status if an echo buffer is corrupted. The ASC should be 3Fh (TARGET OPERATING CONDITIONS HAVE CHANGED) with hopefully a new ASCQ of 0Fh (ECHO BUFFER OVERWRITTEN).

This requires that a target track the initiator that wrote to the echo buffer. A target may implement additional echo buffers to reduce the exceptions.

The READ ECHO BUFFER descriptor option should be required so that the size of the echo buffer can be determined. Sequences longer than 128 or 252 bytes are needed to adequately test the physical layer so the buffer capacity should be modified to match that in table 88. The 252 byte length gives only 63 leading edge transitions, too short to test for some effects.

An implementers note should be added to SPI-3 annex on physical layer integrity checking to deal with initiators that support fairness. Initiators support of fairness during a domain validation sequence may result in an increased number of exception conditions. It is recommended that initiators suspend their fairness algorithm during domain validation.

Changes to SPC-2:

7.15.5 Read Data from echo buffer (1010b)

In this mode the device server transfers data to the application client from the echo buffer. The echo buffer shall transfer the same data as when the WRITE BUFFER command with the mode field set to echo buffer was issued from the same initiator. The BUFFER ID and BUFFER OFFSET fields are ignored in this mode.

~~NOTE 28 A WRITE BUFFER command with the mode field set to echo buffer should be sent prior to the READ BUFFER command; otherwise the READ BUFFER command may terminate with CHECK CONDITION and a sense key of ILLEGAL REQUEST or may return indeterminate data.~~

The READ BUFFER command shall return the same number of bytes of data as received in the prior echo buffer mode WRITE BUFFER command from the same initiator. If a prior echo buffer mode WRITE BUFFER command was not successfully completed the echo buffer mode READ BUFFER command shall terminate in a CHECK CONDITION status, the sense key shall be set to ILLEGAL REQUEST and the additional sense code to COMMAND SEQUENCE ERROR (2Ch/00h). If the data in the echo buffer is overwritten by another initiator the target shall terminate the echo buffer mode READ BUFFER command with a CHECK CONDITION STATUS, the sense key shall be set to ABORTED COMMAND and the additional sense code to ECHO BUFFER OVERWRITTEN (3Fh/0Fh).

The initiator is allowed to send a READ BUFFER command requesting the echo buffer descriptor prior to a WRITE BUFFER command.

If an echo buffer mode WRITE BUFFER command is successfully executed then the initiator is allowed to send multiple echo buffer mode READ BUFFER commands.

7.15.6 Echo buffer descriptor mode (1011b)

In this mode, a maximum of four bytes of READ BUFFER descriptor information is returned. The device server shall return the descriptor information for the echo buffer. If there is no echo buffer implemented, the device server shall return all zeros in the READ BUFFER descriptor. The BUFFER OFFSET field is reserved in this mode. The allocation length should be set to four or greater. The device server shall transfer the lesser of the allocation length or four bytes of READ BUFFER descriptor. The READ BUFFER descriptor is defined as shown in table 90.

Table 90 - Echo Buffer Descriptor

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	Reserved							
1	Reserved							
2	Reserved							
3	Buffer Capacity							

The BUFFER CAPACITY field shall return the size of the echo buffer in bytes aligned to a four-byte boundary.

If the echo buffer is implemented then the echo buffer descriptor shall be implemented.

7.28.8 Write data to echo buffer (1010b)

In this mode the device server transfers data from the application client and stores it in an echo buffer. An echo buffer is assigned in the same manner by the target as it would for a write operation. Data shall be sent aligned on four-byte boundaries. The BUFFER ID and BUFFER OFFSET fields are ignored in this mode.

Note nn: It is recommended that the target assign echo buffers on a per initiator basis to limit the number of exception conditions that may occur in a multi-initiator environment.

Upon successful completion of a WRITE BUFFER command the data shall be preserved in the echo buffer unless there is an intervening command to any logical unit in which case it may be changed.