

Date: August 25, 2000

To: T10 Committee (SCSI)

From: Rob Basham (IBM)

Subject: Explicit State Change Proposal for SSC-2

1. Introduction

1.1 Author Information

Document Owner: Rob Basham

Document Owner Email Address: robbyb@us.ibm.com

Document Owner Phone Number: 520-799-4923

1.2 Change History

1.2.1 Revision 0

- Initial Proposal

1.3 Purpose

This document describes a set of changes to the SCSI Tape command set that makes all state change requests explicit. The command set as currently documented in SSC is such that the meaning of many commands is dependent on preceding commands. For example, the READ command does not contain a logical block number. Which block is read depends on the current position on tape.

The goal of the changes is to allow a device driver to safely redrive commands, perform alternate pathing when a path has a hard failure, perform command queueing where in-order delivery is not guaranteed, etc. without the need for complicated logic to regain positioning agreements between the Initiator and the Target.

2. Implicit State Issues

The following are some of the problems in the current command set that this proposal resolves:

- Relative nature of the SPACE command
- Filemark traversal on READ and SPACE commands
- Implicit logical block number for READ, SPACE, WRITE, WRITE FILEMARKS, VERIFY, RECOVER BUFFERED DATA and ERASE commands.

2.1 Relative Nature of SPACE Command

The SPACE command is problematic in that it asks the drive to space relative to the current position on tape. If for some reason that command failed on the interface, re-driving the command is problematic. Did the SPACE command get to the tape drive or not? On a command timeout at the device driver, there is currently no way to simply re-drive the command because the driver has no idea if the first time around any movement in logical block position took place.

2.2 Filemark Traversal on READ and SPACE

If a filemark is encountered on a READ or SPACE block command, the filemark is traversed and a CHECK CONDITION is reported. This can cause problems for a device driver if the CHECK CONDITION never reaches it for some reason. The drive is performing a state changing action and then telling the device driver about it via Sense Data.

2.3 Implicit Logical Block Number Problem

Many of the commands in the SSC command set have no logical block number associated with them. The drive executes the command based on the current logical block position on tape. If something happens where the Initiator and Target get out of synchronization, there is nothing obvious in the commands themselves that identifies a problem.

3. Proposed Additions

Below are proposed additions to the SSC command set.

3.1 ERASE (10)

The ERASE(10) command contains the IMMED and LONG bits as defined in the current ERASE(6) command. In addition, it contains a (4 or 6?) byte BLOCK ADDRESS field specifying the location at which to start the erase operation.

3.2 READ(10)

The READ(10) command contains TRANSFER LENGTH, SILI and FIXED fields just like the READ(6) command. In addition, it contains a (4 or 6?) byte logical block number specifying where the read operation is to start.

Another difference from the READ(6) command would be that if a filemark is encountered, it is not traversed.

3.3 READ REVERSE (10)

The READ REVERSE(10) command contains TRANSFER LENGTH, BYTORD, SILI and FIXED fields just like the READ REVERSE(6) command. In addition, it contains a (4 or 6?) byte logical block number specifying where the read operation is to start.

3.4 RECOVER BUFFERED DATA(10)

The RECOVER BUFFERED DATA(10) command contains TRANSFER LENGTH, SILI and FIXED fields just like the RECOVER BUFFERED DATA(6) command. In addition, it contains a (4 or 6?) byte logical block number specifying where the operation is to start.

3.5 SPACE(10)

The SPACE command gets two additional fields:

Prohibit File Traversal (PFT) bit - When set to 1, it indicates that the operation is to stay with in the current file. If a filemark is encountered, a CHECK CONDITION status is returned with sense key being set to NO SENSE and the additional sense code qualifier being set to FILEMARK DETECTED. The filemark is not traversed. (An alternative here would be to use a new CODE value).

If the CODE field is not equal to 000b (Space Blocks) and Prohibit File Traversal is set, CHECK CONDITION status shall be returned and the sense key shall be set to ILLEGAL REQUEST and the additional sense code and additional sense code qualifier shall be set to INVALID FIELD IN CDB.

In addition, it contains a (4 or 6?) byte logical block number specifying where the operation is to start.

(It might be worthwhile to move all this function into a LOCATE(10) command instead to make the redrive of a command more direct).

3.6 VERIFY(10)

An additional (4 or 6?) byte logical block number specifying where the operation is to start is added to fields as they exist for the VERIFY(6) command.

3.7 WRITE(10)

An additional (4 or 6?) byte logical block number specifying where the operation is to start is added to fields as they exist for the WRITE(6) command.

3.8 WRITE FILEMARKS(10)

An additional (4 or 6?) byte logical block number specifying where the operation is to start is added to fields as they exist for the WRITE(6) command.

4. Unresolved Issues

- What should the policy be when the block requested isn't current logical block? Some possibilities are to execute the read type commands anyways and respond with a CHECK CONDITION on write type commands. Another possibility is to respond with a CHECK CONDITION in all cases. This author leans towards the first option.
- Do we need any enable/disable mode settings?
- Do we need a partition number?