

Date: October 4, 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

1.2.2 Revision 1

- Command Descriptor Blocks, etc, added

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 re-drive 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 Changes

Below are proposed additions to the SSC-2 Revision 0 document.

The following definitions need to be added to section 5.1 on page 8.

block number- A unique identifier for a logical block. The first logical block on tape is block number 0. If logical block has a number n, the logical block immediately before that block has a block number of n-1 and the logical block immediately after that block has a block number of n+1. Filemarks and setmarks take up a block number.

file number - A unique identifier for a range of blocks ending in a filemark. The first file number is zero. If file number has a value of n, the file number immediately before that a file number of n-1 and the file number immediately after is n+1.

set number - A unique identifier for a range of blocks ending in a setmark. The first set number is zero. If set number has a value of n, the set number immediately before that a set number of n-1 and the set number immediately after is n+1.

The following statement needs to be added to section 5.2.6 Tagged command queueing at te end:

When command queueing, operations that contain explicit state information (usually the logical block number) should be used. Legacy commands from the SSC document should not be used.

After 5.2.8 Direction and Position Information, add the following section:

3.0.1 Matching block number or file number requested with current position

On read type commands, no check for mismatch of current position with block number or partition is necessary. Note that this does not apply to the legacy command set where block numbers were not included.

The write type commands should be checked for a mismatch of current position. If there is a mismatch between the current position and the position specified in the command then a CHECK CONDITION should be returned with sense key ILLEGAL REQUEST and an additional sense code qualifier of SEQUENTIAL POSITIONING ERROR.

Table 1 on page 17 needs to be updated to account for the new commands.

Table 3 on page 20 needs to be updated to account for the new commands.

Table 9 on page 24 needs to be updated to account for the new commands.

Table 10 gets changed to the form below(subclauses and op codes purposely omitted):

Table 1 - Commands for sequential-access devices

Command Name	Flush Write Data	Operation code	Type	Subclause
ERASE (6)	Yes		M	
ERASE (16)	Yes		M	
FORMAT MEDIUM	No		O	
INQUIRY	No		M	
LOAD UNLOAD	Yes		O	
LOCATE (10)	Yes		O	
LOCATE (16)	Yes		O	
LOG SELECT	No		O	
LOG SENSE	No		O	
MODE SELECT (6)	Yes		O	
MODE SELECT (10)	Yes		M	
MODE SENSE (6)	No		O	
MODE SENSE (10)	No		M	
MOVE MEDIUM	Yes		O	
MOVE MEDIUM ATTACHED	Yes		O	
PERSISTENT RESERVE IN	No		O	
PERSISTENT RESERVE OUT	No		O	
PREVENT ALLOW MEDIA REMOVAL	No		O	
READ (6)	Yes		M	
READ (16)	Yes		??	
READ BLOCK LIMITS	No		M	
READ BUFFER	Yes		O	
READ ELEMENT STATUS	No		O	
READ ELEMENT STATUS ATTACHED	No		O	
READ POSITION	No		M	
READ REVERSE (6)	Yes		O	
READ REVERSE (16)	Yes		O	
RECEIVE DIAGNOSTIC RESULTS	No		O	
RECOVER BUFFERED DATA	May		O	
RELEASE(6)	No		M	
RELEASE(10)	No		M	
REPORT DENSITY SUPPORT	No		M	
REPORT LUNS	No		M	
REQUEST SENSE	No		M	
RESERVE(6)	No		M	
RESERVE(10)	No		M	
REWIND	Yes		M	
SEND DIAGNOSTIC	Yes		O	
SET CAPACITY	May		O	
SPACE(6)	May		??	
SPACE(16)	May		O	
TEST UNIT READY	No		O	
VERIFY (6)	Yes		O	
VERIFY (16)	Yes		O	
WRITE (6)	No		??	
WRITE (16)	No		M	
WRITE BUFFER	Yes		O	
WRITE FILEMARKS (6)	May		??	
WRITE FILEMARKS (16)	May		M	

Key: M = Command implementation is mandatory.
O = Command implementation is optional.

On page 27 change the ERASE command to ERASE(6).

3.1 ERASE (12)

The ERASE(12) command ... *add in text here just as it appears in 5.3.1 ERASE Command*

Table 2 - ERASE (12) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED						IMMED	LONG
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10								
11	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION fields indicate the position at which the ERASE command should start. If the current logical position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION fields, a CHECK CONDITION should be returned with sense key ILLEGAL REQUEST and an additional sense code qualifier of SEQUENTIAL POSITIONING ERROR.

3.2 LOCATE (16)

The LOCATE (16) command ... *add in text here just as it appears in 5.3.1 LOCATE (16)*

Command

Table 3 - LOCATE (92) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED				CODE		IMMED	
2	LOGICAL BLOCK ADDRESS							
9								
10	RESERVED							
13								
14	PARTITION							
15	CONTROL							

Add the following text for the code field:

The lcode field is defined in the table below:

Table 4 -

Code	Description	Support
000b	Block	Mandatory
001b	Filemark	Mandatory
010b	Sequential filemarks	Optional
011b	End-of-data	Optional
100b	Setmark	Optional
101b	Sequential setmarks	Optional
110b	BT	Optional
111b	CP	Optional

The text describing the the details of each code come out of the current LOCATE (10) and SPACE (6) command text. The only difference is that on the encountering of a filemark or setmark, the drive does not traverse the mark when reporting up the error. Traversal requires a LOCATE command be sent with the appropriate code (Filemark or Setmark).

On page 32 change the READ command to READ(6).

3.3 READ (16)

The READ(16) command ... *add in text here just as it appears in 5.3.6 READ Command*

Table 5 - READ (16) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED						SILI	FIXED
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10								
11	TRANSFER LENGTH							
13								
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION indicates the position at which the READ command should start. If the current position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION field the drive should locate to that position and return the block as requested.

Also, text should be changed to reflect that filemarks and setmarks are not traversed when encountered.

3.4 READ REVERSE (16)

The READ REVERSE(16) command ... *add in text here just as it appears in 5.3.9 READ*

*REVERSE Command***Table 6 - READ REVERSE (16) command**

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED					BYTORD	SILI	FIXED
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10	LOGICAL BLOCK ADDRESS							
11	TRANSFER LENGTH							
13	TRANSFER LENGTH							
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION indicates the position at which the READ REVERSE command should start. If the current position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION field the drive should locate to that position and return the block as requested.

Also, text should be changed to reflect that filemarks and setmarks are not traversed when encountered.

3.5 RECOVER BUFFERED DATA (16)

The RECOVER BUFFERED DATA(16) command ... *add in text here just as it appears in 5.3.10*

*RECOVER BUFFERED DATA Command***Table 7 - RECOVER BUFFERED DATA(16) command**

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED					SILI		FIXED
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10	LOGICAL BLOCK ADDRESS							
11	TRANSFER LENGTH							
13	TRANSFER LENGTH							
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION indicates the position at which the RECOVER BUFFERED DATA command should start. If the current position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION field the drive should locate to that position and return the block as requested.

Also, text should be changed to reflect that filemarks and setmarks are not traversed when encountered.

3.6 SPACE (16)

If it's not too late, I would propose obsoleting the SPACE(16) command and requiring use of the LOCATE (16) command.

3.7 VERIFY (16)

The VERIFY(16) command ... *add in text here just as it appears in 5.3.16 VERIFY Command*

Table 8 - VERIFY (16) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED				IMMED	BYTCMP	FIXED	
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10	LOGICAL BLOCK ADDRESS							
11	VERIFICATION LENGTH							
13	VERIFICATION LENGTH							
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION indicates the position at which the VERIFY command should start. If the current position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION field the drive should locate to that position and return the block as requested.

Also, text should be changed to reflect that filemarks and setmarks are not traversed when encountered.

3.8 WRITE (16)

The WRITE (16) command ... *add in text here just as it appears in 5.3.17 WRITE Command*

Table 9 - WRITE (16) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED							FIXED
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10	LOGICAL BLOCK ADDRESS							
11	TRANSFER LENGTH							
13	TRANSFER LENGTH							
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION fields indicate the position at which the WRITE command should start. If the current logical position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION fields, a CHECK CONDITION should be returned with sense key ILLEGAL REQUEST and an additional sense code qualifier of SEQUENTIAL POSITIONING ERROR.

3.9 WRITE FILEMARKS (16)

The WRITE FILEMARKS (16) command ... *add in text here just as it appears in 5.3.17 WRITE*

*FILEMARKS Command***Table 10 - WRITE FILEMARKS (16) command**

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE							
1	RESERVED					WSMK	FIXED	
2	PARTITION							
3	LOGICAL BLOCK ADDRESS							
10	LOGICAL BLOCK ADDRESS							
11	TRANSFER LENGTH							
13	TRANSFER LENGTH							
14	RESERVED							
15	CONTROL							

Add the following text for the logical block number and partition fields:

The LOGICAL BLOCK ADDRESS and PARTITION fields indicate the position at which the WRITE FILEMARKS command should start. If the current logical position does not match what is in the LOGICAL BLOCK ADDRESS and PARTITION fields, a CHECK CONDITION should be returned with sense key ILLEGAL REQUEST and an additional sense code qualifier of SEQUENTIAL POSITIONING ERROR.