

To: T10 Membership
From: Paul Suhler, Seagate Technology
Date: 29 February 2000
Subject: Support for Large Block Addresses in SSC-2
Document: T10/00-135r1



1 Changes from Previous Versions

1.1 Changes in 00-135r1

- Changed assumption 4 to transfer lengths less than 2^{24} . Dropped changes to all commands other than LOCATE, READ POSITION, and SPACE.
- Modified table formats to match those in SSC source.
- Use operation codes 91h and 92h for SPACE (16) and LOCATE (16), respectively, as suggested by Ralph Weber.
- Snarfed and modified entire READ POSITION command clause to show changes in context.
- Changed separate bit fields in READ POSITION CDB byte 1 to a five-bit SERVICE ACTION field, and changed bytes 7 and 8 to an ALLOCATION LENGTH field, as suggested by Paul Entzel.
- Changed bytes 2 and 3 in the proposed extended data form to an ADDITIONAL LENGTH field.

2 Background

There now exist tape drives whose media are large enough that they could hold more blocks than could be addressed by existing SSC commands. For example, first-generation LTO drives have cartridge capacities of 200 GB, and the LTO roadmap has 1.6 TB media in generation four. Even with only 200 GB, a tape with small block sizes could hold more blocks than the existing LOCATE command can address with its four-byte BLOCK ADDRESS field. A 1.6 TB medium would max out the logical block address with block sizes of less than 400 bytes.

While workarounds will be possible without command set changes, we should use this opportunity to address the command set issues before we have to invent vendor-unique solutions which will become legacy support problems.

The January 2000 working group ranked general solutions by desirability:

1. New 16-byte commands
2. New service actions under the variable length CDB
3. New 12-byte commands
4. A new bit in the current commands

This proposal addresses the first solution, sixteen-byte commands.

3 Assumptions

The following assumptions were derived by looking at the command and data fields for SSC commands to see where fields may overflow. Assumptions 1 through 3 argue for changes in block address fields to eight; the rest argue for no changes in other fields. These assumptions should be considered carefully and discussed.

1. Number of blocks on a medium may be more than 2^{32} .
2. Absolute block addressing may specify more than 2^{32} blocks.
3. Relative block addressing may specify more than 2^{32} blocks.
4. Read and write transfer lengths will be less than 2^{24} .
5. Number of blocks in a device buffer will be less than 2^{24} .
6. Number of bytes in a device buffer will be less than 2^{24} .
7. Number of setmarks or filemarks in a partition will be less than 2^{32} .
8. Number of partitions on a medium will be less than 2^8 .
9. Number of blocks in a file will be less than 2^{24} .

The February 2000 working group concurred with revising assumption 4, i.e., that transfers of more than 16 MB do not need support at this time.

4 Command Changes

4.1 Summary

Command Name	Operation Code	Type	SSC Clause	Comment
ERASE	19h	M	5.3.1	No change
FORMAT MEDIUM	04h	O	5.3.2	No change
LOAD UNLOAD	1Bh	O	5.3.3	No change
LOCATE	2Bh	O	5.3.4	Use CDB format from 4.2 for Logical Block Address
READ	08h	M	5.3.5	No change
READ BLOCK LIMITS	05h	M	5.3.6	No change
READ POSITION	34h	M	5.3.7	Use CDB format with Service Action and Allocation Length and new extended data format in 4.3.
READ REVERSE	0Fh	O	5.3.8	No change
RECOVER BUFFERED DATA	14h	O	5.3.9	No change
REPORT DENSITY SUPPORT	44h	M	5.3.10	No change
REWIND	01h	M	5.3.11	No change
SPACE	11h	M	5.3.12	Use CDB format from 4.4 for Count
VERIFY	13h	O	5.3.13	No change
WRITE	0Ah	M	5.3.14	No change
WRITE FILEMARKS	10h	M	5.3.15	No change

4.2 LOCATE (16) Command

4.2.1 Discussion

Assumption 2, the need for addressing more than 2^{32} blocks on a medium, justifies changing the LOCATE command.

4.2.2 Proposed Changes

Change all references to the original LOCATE command to "LOCATE (10)."

Add the following clause:

5.3.X LOCATE (16) Command

The LOCATE (16) command (see Table 1) operates identically to the LOCATE (10) command, but allows specifying a Logical Block Address up to eight bytes in length.

Table 1 – LOCATE (16) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (92h)								
1		Reserved					BT	CP	IMMED	
2	(MSB)	LOGICAL BLOCK ADDRESS								
3										
4										
5										
6										
7										
8										
9										
10		Reserved								
11		Reserved								
12		Reserved								
13		Reserved								
14		PARTITION								
15		CONTROL								

4.3 READ POSITION Command

4.3.1 Discussion

The difficulty with the READ POSITION command is that the short form of the data uses four-byte fields for FIRST BLOCK LOCATION and LAST BLOCK LOCATION. (The long form of the data has an eight-byte field for BLOCK NUMBER, and thus requires no change.) However, because the short form provides the number of blocks and bytes in the device buffer in addition to the block locations, it is useful and should be adapted.

Therefore, we propose a new extended version of the short form of data providing eight-byte fields for FIRST BLOCK LOCATION and LAST BLOCK LOCATION. The fields appear in a new order to allow eight-byte alignment of the eight-byte fields.

A READ POSITION command specifying the short data form when there are 2^{32} or more blocks on the medium would result in the PERR bit's being set; therefore no new operation is needed for this case.

4.3.2 Proposed Changes

Change clause 5.3.7 as follows to (i) incorporate the TCLP, LONG, and BT fields into a new SERVICE ACTION field, and (ii) add an ALLOCATION LENGTH field that will be meaningful for service actions numbered 08h and higher. Red text is new or changed. Yellow highlighted text will require hot links to other clauses.

5.3.7 READ POSITION command

The READ POSITION command (see Table 2) reports the current position and provides information about any data blocks, filemarks and/or setmarks in the buffer. No medium movement shall occur as a result of responding to the command.

Table 2 — READ POSITION command

Bit Byte	7	6	5	4	3	2	1	0	
0	OPERATION CODE (34h)								
1	Reserved			SERVICE ACTION					
2	Reserved								
3	Reserved								
4	Reserved								
5	Reserved								
6	Reserved								
7	(MSB)	ALLOCATION LENGTH							
8								(LSB)	
9	CONTROL								

The service actions defined for the READ POSITION command are shown in Table 3.

Table 3 — READ POSITION service action codes

Code	Name	Description	Implementation Requirements
------	------	-------------	-----------------------------

00h	SHORT FORM -- BLOCK ID	Device server shall return 20 bytes of data with the FIRST BLOCK LOCATION, the LAST BLOCK LOCATION, and BLOCK NUMBER fields as block identifier values (see 5.2.7), (relative to a partition).	Mandatory
01h	SHORT FORM -- VENDOR-SPECIFIC	Device server shall return 20 bytes of data with the FIRST BLOCK LOCATION, the LAST BLOCK LOCATION, and BLOCK NUMBER fields as vendor-specific values.	Optional
02h	Reserved	Illegal request	
03h	Reserved	Illegal request	
04h	Reserved	Illegal request	
05h	Reserved	Illegal request	
06h	LONG FORM	Device server shall return 32 bytes of data.	Optional
07h	Reserved	Illegal request	
08h	EXTENDED FORM	Device server shall return 28 bytes of data up to the maximum length specified by the ALLOCATION LENGTH field.	Optional
09h - 1Fh	Reserved		

If the device server does not implement the specified service action code, then the command shall be terminated with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code and an additional sense code qualifier set to INVALID FIELD IN CDB.

To maintain compatibility with earlier implementations, the service action codes 02h, 03h, 04h, 05h, and 07h shall not be implemented.

For service action codes of 00h, 01h, and 06h, the ALLOCATION LENGTH field in the CDB shall be zero. If it is not, then the command shall be terminated with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code and an additional sense code qualifier set to INVALID FIELD IN CDB.

For service action codes of 08h and higher, the ALLOCATION LENGTH field in the CDB indicates how much space has been reserved for the returned parameter list. If the length is not sufficient to contain the entire parameter list, the first portion of the list shall be returned. This shall not be considered an error. If the remainder of the list is required, the application client should send a new READ POSITION command with a ALLOCATION LENGTH field large enough to contain the entire parameter list.

Table 4 indicates the READ POSITION data that shall be returned if the SERVICE ACTION field is 00h or 01h.

Table 4 — READ POSITION data format, short form

Bit Byte	7	6	5	4	3	2	1	0
0	BOP	EOP	BCU	BYCU	Rsvd	BPU	PERR	Rsvd
1	PARTITION NUMBER							
2	Reserved							
3	Reserved							
4	(MSB)							
:	FIRST BLOCK LOCATION							
7							(LSB)	
8	(MSB)							
:	LAST BLOCK LOCATION							
11							(LSB)	
12	Reserved							
13	(MSB)							
:	NUMBER OF BLOCKS IN BUFFER							
15							(LSB)	
16	(MSB)							
:	NUMBER OF BYTES IN BUFFER							
19							(LSB)	

A beginning-of-partition (BOP) bit of one indicates that the logical unit is at the beginning-of-partition in the current partition. A BOP bit of zero indicates that the current logical position is not at the beginning-of-partition.

An end-of-partition (EOP) bit of one indicates that the logical unit is positioned between early-warning and end-of-partition in the current partition. An EOP bit of zero indicates that the current logical position is not between early-warning and end-of-partition.

A block count unknown (BCU) bit of one indicates that the NUMBER OF BLOCKS IN BUFFER field does not represent the actual number of blocks in the buffer. A BCU bit of zero indicates that the NUMBER OF BLOCKS IN BUFFER field is valid.

A byte count unknown (BYCU) bit of one indicates that the NUMBER OF BYTES IN BUFFER field does not represent the actual number of bytes in the buffer. A BYCU bit of zero indicates that the NUMBER OF BYTES IN BUFFER field is valid.

A block position unknown (BPU) bit of one indicates that the first and last block locations are not currently known or not otherwise obtainable. A BPU bit of zero indicates that the FIRST BLOCK LOCATION and LAST BLOCK LOCATION fields contain valid position information.

A position error (PERR) bit of one indicates that the logical unit is unable to report the correct position due to an overflow of any of the returned position data. A PERR bit of zero indicates that an overflow has not occurred in any of the returned position data fields.

The PARTITION NUMBER field reports the partition number for the current logical position. If the logical unit only supports one partition for the medium, this field shall be set to zero.

The FIRST BLOCK LOCATION field indicates the block identifier associated with the current logical position. The value shall indicate the block identifier of the next data block to be transferred between an application client and the device server if a READ or WRITE command is issued.

The LAST BLOCK LOCATION field indicates the block identifier (see 5.2.7) associated with the next block to be transferred from the buffer to the medium. The value shall indicate the block identifier of the next data block to be transferred between the buffer and the medium. If the buffer does not contain a whole block of data or is empty, the value reported for the last block location shall be equal to the value reported for the first block location.

NOTE 1 The information provided by the FIRST BLOCK LOCATION and LAST BLOCK LOCATION fields may be used in conjunction with the LOCATE command to position the medium at the appropriate logical block on another device in the case of unrecoverable errors on the first device.

The NUMBER OF BLOCKS IN BUFFER field indicates the number of data blocks in the buffer of the logical unit that have not been written to the medium.

The NUMBER OF BYTES IN BUFFER field indicates the total number of data bytes in the buffer of the logical unit that have not been written to the medium.

Table 5 indicates the format of the READ POSITION data that shall be returned if the SERVICE ACTION field is 06h.

Table 5 — READ POSITION data format, long form

Bit	7	6	5	4	3	2	1	0	
Byte									
0	BOP	EOP	Reserved		MPU	BPU	Reserved		
1	Reserved								
2	Reserved								
3	Reserved								
4	(MSB)								
:	PARTITION NUMBER								
7									(LSB)
8	(MSB)								
:	BLOCK NUMBER								
15									(LSB)
16	(MSB)								
:	FILE NUMBER								
23									(LSB)
24	(MSB)								
:	SET NUMBER								
31									(LSB)

The BOP, EOP, and PARTITION NUMBER fields are as defined in the READ POSITION data returned when the TCLP bit is set to zero.

A block position unknown (BPU) bit of one indicates that the partition number or block number are not known or accurate reporting is not currently available. A BPU bit of zero indicates that the PARTITION NUMBER and BLOCK NUMBER fields contain valid position information.

A mark position unknown (MPU) bit of one indicates the file number and set number are not known or accurate reporting is not currently available. A MPU bit of zero indicates the FILE NUMBER and SET NUMBER fields contain valid position information.

The MPU and BPU bits shall be set to one only if the logical unit is unable to accurately assume or does not know the current mark or block, respectively.

The PARTITION NUMBER field reports the partition number for the current logical position. If the logical unit only supports one partition for the medium, this field shall be set to zero.

The BLOCK NUMBER shall report the number of logical blocks between beginning-of-partition and the current logical position. Setmarks and filemarks count as one logical block each.

The FILE NUMBER shall report the number of filemarks between beginning-of-partition and the current logical position.

The SET NUMBER shall report the number of setmarks between beginning-of-partition and the current logical position.

NOTE 2 The reported SET NUMBER value is not affected by the value of the RSMK bit in the device configuration page.

Table 6 indicates the format of the READ POSITION data that shall be returned if the SERVICE ACTION field is 08h.

Table 6 – READ POSITION data format, extended form

Byte	Bit	7	6	5	4	3	2	1	0
0		BOP	EOP	BCU	BYCU	Rsvd	BPU	PERR	Rsvd
1		PARTITION NUMBER							
2	(MSB)	ADDITIONAL LENGTH (18h)							
3									
4		Reserved							
5	(MSB)								
6		NUMBER OF BLOCKS IN BUFFER							
7									
8	(MSB)								
...		FIRST BLOCK LOCATION							
15									
16	(MSB)								
...		SECOND BLOCK IN BUFFER							
23									
24	(MSB)								
...		NUMBER OF BYTES IN BUFFER							
27									

The fields are defined the same as for the corresponding fields in the short data form (Table 4).

The ADDITIONAL LENGTH field shall contain 18h. If the information transferred to the Data-In Buffer is truncated because of an insufficient allocation length value, the ADDITIONAL LENGTH field shall not be altered to reflect the truncation.

[Note: The above behavior of the ADDITIONAL LENGTH field is standard according to SPC-2, clause 4.2.5. Is it desirable or should the value be adjusted when truncation occurs?]

4.4 SPACE (16) Command

4.4.1 Discussion

Assumption 3, the need for relative addressing of more than 2³² blocks, justifies changing the SPACE command. Without this change, a workaround would be to issue multiple SPACE commands, since the block count is relative to the current position. However, on some media formats, such as serpentine, this could mean changing directions between commands, which would be more time-consuming. The larger block count would allow the device to move the medium directly to the desired location, without the intermediate stops.

4.4.2 Proposed Changes

Change all references to the original SPACE command to “SPACE (6).”

Add the following clause:

5.3.X SPACE (16) Command

The SPACE (16) command (see Table 7) operates identically to the SPACE (6) command, but allows specifying a count up to eight bytes in length.

Table 7 – SPACE (16) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (91h)								
1		Reserved					CODE			
2	(MSB)									
3										
4										
5										
6		COUNT								
7										
8										
9		(LSB)								
10		Reserved								
11		Reserved								
12		Reserved								
13		Reserved								
14		Reserved								
15		CONTROL								