

Date: July 31, 2003

To: T10 Committee (SCSI)

From: George Penokie (IBM/Tivoli)

Subject: End-to-End Data Protection

1 Overview

There is an need (real or imagined) for a standardized end-to-end data protection mechanism to be defined. The logical place to such a definition is the SCSI command and architecture standards, as most storage uses SCSI commands to read/write data to and from storage devices. What follows is a proposal that provides a set of SCSI tools that will enable end-to-end data protection. This set of SCSI tools are defined to accomplish this goal:

- a) without interfering with existing proprietary methods;
- b) with a minimum of options; and
- c) by defining minimal changes to CDBs while maintaining backward compatibility.

The set of SCSI tools will consist of the following:

- a) Two level data protection on each data block transferred across the interconnect that consists of;
 - A) A 4-byte CRC that covers the user data of the data block. The CRC is generated at or before the application client and preserved at the logical unit.
 - B) A 4-byte incrementing LBA tag. The incrementing LBA tag is set by the application client during write operation to the value of the least significant 4 bytes of the write command's LBA field on the first data block transferred and incremented by one on each data block transferred until all the blocks for the command have been transferred. The increment LBA tag values for each data block that is read back shall be the same value that was received for that data block.
- b) A bit in the non-Read Read CDBs (e.g. VERIFY) to allow a logical unit to return the protection data.
 - A) If zero then do not transmit any protection information. The logical unit shall not check the contents of the protection fields.
 - B) If one then transmit the protection information. The logical unit may determine if the data block is valid by checking the contents of the protection fields. If the logical unit determines there is a an error it shall generate a check condition. A read to a logical unit that has not been formatted to transmit the data protection fields may fail with a check condition. In the case where the logical unit has not been formatted and does not check the bit the contents of the protection information is unpredictable and as a result should cause an error at the application client.
- c) A two bit field in the READ commands (excluding the READ (6) command) that would control the reading and checking of protection data.
 - A) If 00b then do not transmit any protection information. ~~The logical unit shall may not check the contents of the protection fields~~ If the logical unit has been formatted with protection information the logical unit may determine if the data block is valid by checking the contents of the protection fields. If the logical unit determines there is a an error it shall generate a check condition.
 - B) If 01b then transmit the protection information. The logical unit may determine if the data block is valid by checking the contents of the protection fields. If the logical unit determines there is a an error it shall generate a check condition. A read to a logical unit that has not been formatted to transmit the data protection fields may fail with a check condition. In the case where the logical unit has not been formatted and does not check the bit the contents of the protection information is unpredictable and as a result should cause an error at the application client.
 - C) If 11b then transmit the protection information. The logical unit shall not check the contents of the protection fields. A read to a logical unit that has not been formatted to transmit the data protection fields may fail with a check condition. In the case where the logical unit has not been formatted and does not check the bit the contents of the protection information is unpredictable and as a result should cause an error at the application client.
- d) A bit in Write CDBs to allow the protection information to be written with no checks.
 - A) If 00h then preserve the contents of the protection fields (e.g., write to media, store in non-volatile memory, recalculate on read back). The logical unit shall determine if the data block is valid by

checking the contents of the protection fields. If the logical unit determines there is an error it shall generate a check condition. If the logical unit has not been formatted to accept protection information it shall generate a check condition.

- B) If 01h then preserve the contents of the protection fields (e.g., write to media, store in non-volatile memory, recalculate on read back). The logical unit shall not check the contents of the protection fields. If the logical unit has not been formatted to accept protection information it shall generate a check condition.
- C) If 10h then the contents of the protection fields shall not be preserved. The logical unit shall determine if the data block is valid by checking the contents of the protection fields. If the logical unit determines there is an error it shall generate a check condition. In the case where the logical unit has not been formatted with protection fields and does not check the CDB protection field the logical unit's response to the command is unpredictable.
- e) A bit in the Format CDB to cause 8 bytes to be added to the block size of the logical unit being formatted.
 - A) If zero then format the medium to the block length defined in the mode parameter block descriptor of the Mode parameter header.
 - B) If one then format the medium to the block length defined in the mode parameter block descriptor of the Mode parameter header plus 8 (e.g., if block length = 512 the formatted block length is 520). The block length shall be a multiple of four. If the block length is not a multiple of four the logical unit shall generate a check condition.
- f) All commands that request block length information (e.g., Read Capacity, Mode Sense) shall return the block size of the data excluding the eight bytes of protection information (e.g., a 520 byte data block on a device formatted with protection information returns 512 in the block length field).
- g) A two bit field in the Standard Inquiry Data to indicate support of protection data.
 - A) If 00b then no protection is supported.
 - B) If 10b then protection is supported but not enabled
 - C) If 11b then protection is supported and enabled.
- h) A bit in a mode page that forces the logical unit to write to media the contents of the crc field.
 - A) If zero then the contents of the crc field shall be preserved and may be written to media.
 - B) If one then the contents of the crc field shall be written to media.

SBC-2 additions

4.x Data protection model

4.0.1 Data protection overview

This data protection model provides for the protection of the data while it is being transferred between a sender and a receiver. The protection data is generated at the application layer and may be checked by any object along the I_T_L nexus. Once received, the protection information shall be retained (e.g., write to media, store in non-volatile memory, recalculate on read back) by the device server until overwritten (e.g., power loss and reset events have no effect on the retention of protection information).

4.0.2 Protected data

The protection data consists of two fields appended to each block of data that contains:

- a) a cyclic redundancy check (CRC); and
- b) a logical block address tag.

See figure 1 for the placement of the CRC and LOGICAL ADDRESS TAG fields.

Table 1 — Protected data block format

Byte\Bit	7	6	5	4	3	2	1	0
0	DATA BLOCK							
n								
n + 1	(MSB)	CRC						(LSB)
n + 4								
n + 5	(MSB)	LOGICAL ADDRESS TAG						(LSB)
n + 8								

The data block shall contain user data.

The crc field contains the CRC (see 4.0.3) of the contents of the DATA BLOCK field.

The LOGICAL ADDRESS TAG field is set to the least significant four bytes of the logical block address to which the data block is to be associated with. The first data block transmitted shall contain the least significant four bytes of the logical block address contained in the LOGICAL BLOCK ADDRESS field of the command associated with the data being transferred. Each subsequent data block's LOGICAL ADDRESS TAG field shall contain the logical address tag of the previous data block plus one.

4.0.3 CRC protection

If data protection is enabled, the application client shall append a CRC to each block of data to be transmitted. The CRC shall be generated from the contents of the DATA BLOCK field.

Annex C contains information on CRC generation/checker implementation.

Table 2 defines the CRC polynomials.

Table 2 — CRC polynomials

Function	Definition
F(x)	A polynomial of degree k-1 that is used to represent the k bits of the data block covered by the CRC. For the purposes of the CRC, the coefficient of the highest order term shall be the MSB of the DATA BLOCK field.
L(x)	A degree 31 polynomial with all of the coefficients set to one: $L(x) = x^{31} + x^{30} + x^{29} + \dots + x^2 + x^1 + 1$ (i.e., L(x) = FFFFFFFFh)
G(x)	The standard generator polynomial: $G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$ (i.e., G(x) = 1_04C11DB7h)
R(x)	The remainder polynomial, which is of degree less than 32.
P(x)	The remainder polynomial on the receive checking side, which is of degree less than 32.
Q(x)	The greatest multiple of G(x) in $(x^{32} \times F(x)) + (x^k \times L(x))$
Q'(x)	$x^{32} \times Q(x)$
M(x)	The sequence that is transmitted.
M'(x)	The sequence that is received.
C(x)	A unique polynomial remainder produced by the receiver upon reception of an error free sequence. This polynomial has the value: $C(x) = x^{32} \times \frac{L(x)}{G(x)}$ $C(x) = x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^8 + x^6 + x^5 + x^4 + x^3 + x + 1$ (i.e., C(x) = C704DD7Bh)

4.0.4 CRC generation

The equations that are used to generate the CRC from F(x) are as follows. All arithmetic is modulo 2.

$$\text{CRC value in data block} = L(x) + R(x) = \text{one's complement of } R(x)$$

NOTE 1 - Adding L(x) (all ones) to R(x) produces the one's complement of R(x); this equation is specifying that the R(x) is inverted before it is transmitted.

The CRC is calculated by the following equation:

$$\frac{(x^{32} \times F(x)) + (x^k \times L(x))}{G(x)} = Q(x) + \frac{R(x)}{G(x)}$$

The following equation specifies that the CRC is appended to the end of F(x):

$$M(x) = x^{32} \times F(x) + \text{CRC}$$

The bit order of F(x) presented to the CRC function is MSB to LSB four bytes at a time until the contents of the DATA BLOCK field are all processed. This order is shown in figure 1

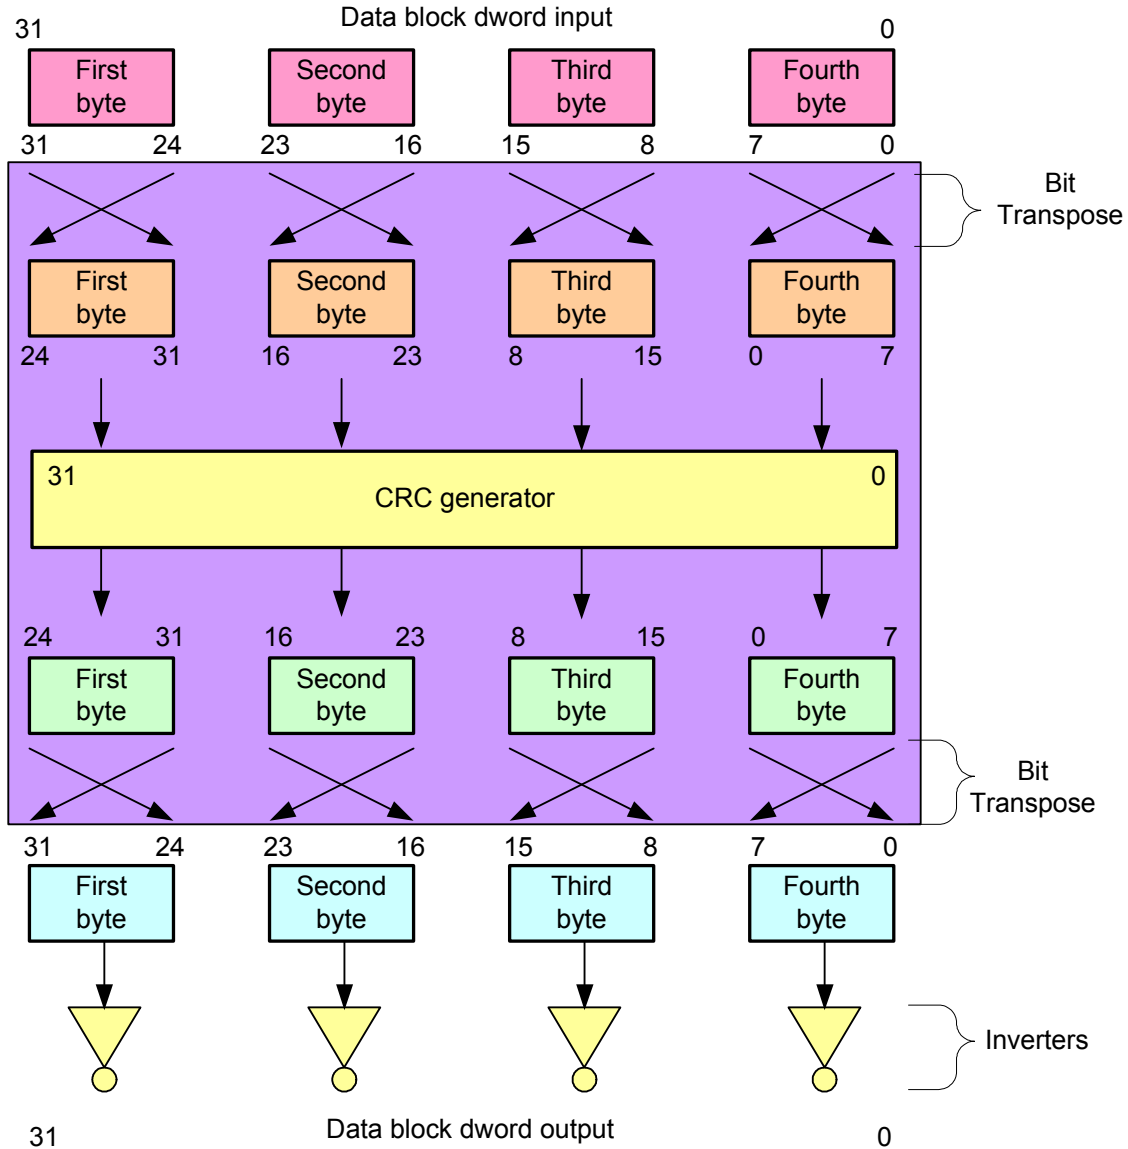


Figure 1 — CRC generator bit order

4.0.5 CRC checking

The received sequence $M'(x)$ may differ from the transmitted sequence $M(x)$ if there are transmission errors. The process of checking the sequence for validity involves dividing the received sequence by $G(x)$ and testing the remainder. Direct division, however, does not yield a unique remainder because of the possibility of leading zeros. Thus a term $L(x)$ is prepended to $M'(x)$ before it is divided. Mathematically, the received checking is shown by the following equation:

$$x^{32} \times \frac{M'(x) + (x^K \times L(x))}{G(x)} = Q'(x) + \frac{P(x)}{G(x)}$$

In the absence of errors, the unique remainder is the remainder of the division as shown by the following equation:

$$\frac{P(x)}{G(x)} = x^{32} \times \frac{L(x)}{G(x)} = C(x)$$

The bit order of F(x) presented to the CRC checking function is the same order as the CRC generation bit order (see figure 1).

4.0.6 Application of protected data

Before an application client transmits or receives protected data it shall:

- 1) Determine if a logical unit supports protected data using the INQUIRY command (see x.x.x);
- 2) If protected data is supported then determine if the logical unit is formatted to accept protected information using the INQUIRY command (see x.x.x);
- 3) If the logical unit supports protected information and is not formatted to accept protected information then format the logical unit with protected information usage enabled;
- 4) If the logical unit supports protected information and is formatted to accept protected information then the application client may use read commands that support protected information and shall use write commands that support protected information.

4.0.7 Protected data commands

The enabling of protection information enables fields in some commands that instruct the device server on the handling of the protection information. The detailed description of each commands protection information fields are defined in the individual command descriptions.

The commands that are affected when protection information is enabled are:

- a) EXTENDED COPY (See SPC-3);
- b) FORMAT UNIT; (Editing Note: Byte 1, Bit 7 - If the Initialization Pattern is used then the target needs to generate valid protection information)
- c) INQUIRY (See SPC-3);
- d) ~~PRE FETCH (10)/(16); (Editing Note: Byte 1, Bit 7);~~
- e) READ ~~(6)~~(10)/(12)/(16); (Editing Note: Byte 1, Bits 7 and 6 except for READ (6))
- f) READ LONG; ~~(Editing Note: Byte 1, Bits 7 and 6)~~
- g) REASSIGN BLOCKS; (Editing Note: Need comment in command description that the reassigned block needs valid protection information written with it.)
- h) REBUILD (16)/(32); ~~(Editing Note: Byte 1, Bit 7 on (16) and Byte 10, Bit 7 on (32) — The logical unit needs to generate protection information on the rebuilt data)~~
- i) REGENERATE (16)/(32); ~~(Editing Note: Byte 1, Bit 7 on (16) and Byte 10, Bit 7 on (32))~~
- j) SYNCHRONIZE CACHE (10)/(16); (Editing Note: Need comment in command description that the blocks that are synchronized need valid protection information written with it.)
- k) VERIFY (10)/(12)/(16); (Editing Note: Need comment in command description that when blocks are verified the logical unit may use protection information as part of the verification. If it fails because of the protection information then a new ASCQ needs to be returned.)
- l) WRITE ~~(6)~~(10)/(12)/(16); (Editing Note: Byte 1, Bit 7 and 6 except for WRITE (6))
- m) WRITE AND VERIFY (10)/(12)/(16); (Editing Note: Byte 1, Bit 7 and 6, also need comment in command description that when blocks are verified the logical unit may use protection information as part of the verification. If it fails because of the protection information then a new ASCQ needs to be returned.)
- n) WRITE LONG; (Editing Note: Need comment in command description that the data to be written needs to include protection information.)
- o) WRITE SAME (10)/(12); (Editing Note: Need comment in command description that as the blocks of data are written the logical unit needs to generate and write valid protection information.)
- p) XDREAD (10)/(32); (Editing Note: Byte 1, Bit 7 and 6 on (16) and Byte 10, Bit 7 and 6 on (32))
- q) XDWRITE (10)/(32); (Editing Note: Byte 1, Bit 7 and 6 on (16) and Byte 10, Bit 7 and 6 on (32))
- r) XDWRITE EXTENDED (16)/(32)/(64); (Editing Note: Byte 1, Bit 7 and 6 on (16) and Byte 10, Bit 7 and 6 on (32) and (64))
- s) XDWRITEREAD (10)/(32); and (Editing Note: For the Write control Byte 1, Bit 7 and 6 on (16) and Byte 10, Bit 7 and 6 on (32) and (64). For the Read control Byte 1, Bit 6 on (16) and Byte 10, Bit 6 on (32) and (64))
- t) XPWRITE (10)/(32). (Editing Note: Byte 1, Bit 7 and 6 on (16) and Byte 10, Bit 7 and 6 on (32))

If a WRITE (6) command is received after protection information is enabled the device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID COMMAND OPERATION CODE.

A READ (6) command may be sent to a logical unit that has protection information enabled but there is no option available to transmit the protection information.

Commands that result in the return of the length in bytes of each logical block (e.g., MODE SENSE, READ CAPACITY) shall return the length of the user data and shall not include the length of the protection information (e.g., if the user data plus the protection information is equal to 520 byte the length in bytes returned is 512 bytes).

4.0.8 FORMAT UNIT command

4.0.8.1 FORMAT UNIT command overview

The FORMAT UNIT command (see table 3) formats the medium into application client addressable logical blocks per the application client defined options. In addition, the medium may be certified and control structures may be created for the management of the medium and defects. The degree that the medium is altered by this command is vendor-specific.

Table 3 — FORMAT UNIT command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (04h)							
1	FMTPDATA	Reserved	ONGLIST	FMTDATA	CMPLIST	DEFECT LIST FORMAT		
2	Vendor specific							
3	INTERLEAVE							
4								
5	CONTROL							

[An format protection data \(FMTPDATA\) bit of zero indicates to format the medium to the block length specified in the mode parameter block descriptor of the Mode parameter header \(see x.x.x\). A FMTPDATA bit of one indicates to format the medium block length specified in the mode parameter block descriptor of the Mode parameter header plus eight \(e.g., if the block length equals 512 the formatted block length is 520\).](#)

...

An initialization pattern (IP) bit of zero indicates that an initialization pattern descriptor is not included and that the device server shall use its default initialization pattern. An IP bit of one indicates that an initialization pattern descriptor (see 5.2.2.3) is included in the FORMAT UNIT parameter list immediately following the defect list header. [If the FMTPDATA bit and the IP bit are both set to one then valid protection data shall be written to the medium for each logical block.](#)

4.0.9 READ (10) command

The READ (10) command (see table 4) requests that the device server transfer data to the application client. The most recent data value written in the addressed logical block shall be returned.

Table 4 — READ (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (28h)							
1	RDPROTECT	Reserved	DPO	FUA	Reserved		RELADR	
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	Reserved							
7	(MSB) _____							
8	TRANSFER LENGTH _____ (LSB)							
9	CONTROL							

The [RDPROTECT](#) field is defined in table 5.

Table 5 — RDPROTECT field

Value	Description
00b	Do not transmit any protection information. If the logical unit has been formatted with protection information the logical unit may determine if the data block is valid by checking the contents of the protection fields. If the logical unit determines there is an error the command shall be terminated with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to UNRECOVERED READ ERROR.
01b	Transmit the protection information. The logical unit may determine if the data block is valid by checking the protection information. If the logical unit determines there is an error the command shall be terminated with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to UNRECOVERED READ ERROR. A read command to a logical unit that has not been formatted with protection information may fail with a CHECK CONDITION status. If so the sense data shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN CDB. If the logical unit has not been formatted with protection information and does not check the PROTECT field then the contents of the protection information is unpredictable.
10b	Reserved
11b	Transmit the protection information. The logical unit shall not check the contents of the protection fields. A read command to a logical unit that has not been formatted with protection information may fail with a CHECK CONDITION status. If a check condition occurs the sense data shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN CDB. If the logical unit has not been formatted with protection information and does not check the RDPROTECT field then the contents of the protection information is unpredictable.

4.0.10 READ (12) command

The READ (12) command (see table 6) requests that the device server transfer data to the application client from the medium.

Table 6 — READ (12) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (A8h)							
1	RDPROTECT	Reserved	DPO	FUA	Reserved		RELADR	
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	(MSB) _____							
9	TRANSFER LENGTH _____ (LSB)							
10	Reserved							
11	CONTROL							

4.0.11 READ (16) command

The READ (16) command (see table 7) requests that the device server transfer data to the application client.

Table 7 — READ (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (88h)							
1	RDPROTECT		Reserved	DPO	FUA	Reserved		RELADR
2	(MSB) _____							
9	LOGICAL BLOCK ADDRESS _____ (LSB)							
10	(MSB) _____							
13	TRANSFER LENGTH _____ (LSB)							
14	Reserved							
15	CONTROL							

4.0.12 READ LONG command

The READ LONG command (see table 41) requests that the device server transfer data to the application client. The data passed during the READ LONG command is vendor-specific, but shall include the data bytes, [any protection information](#), and the ECC bytes recorded on the medium. The most recent data written, or to be written, in the addressed logical block shall be returned. READ LONG is independent of the Read-Write Error Recovery mode page but does allow retries.

4.0.13 REASSIGN BLOCKS command

The REASSIGN BLOCKS command (see table 42) requests the device server to reassign the defective logical blocks [and that logical blocks protection information](#) to another area on the medium set aside for this purpose. The device server should also record the location of the defective logical blocks to the grown defect list if such a list is supported. More than one physical or logical block may be relocated by each defect descriptor sent by the application client. This command does not alter the contents of the PLIST (see 4.0.8).

4.0.14 REBUILD (16) Command

The REBUILD (16) command (see table 46) requests that the target write to the medium the XOR data generated from the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data [and protection information, if any](#). READ (10) should be used for accesses to SCSI devices supporting less than 2 Terabytes, and READ (16) should be used for accesses to SCSI devices supporting greater than or equal to 2 Terabytes

4.0.15 REBUILD (32) Command

The REBUILD (32) command (see table 50) requests that the target write to the medium the XOR data generated from the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data [and protection information, if any](#).

4.0.16 REGENERATE (16) command

The REGENERATE (16) command (see table 53) requests that the target write to the buffer the XOR data generated from its own medium and the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data [and protection information, if any](#). The resulting XOR data is retained in the target's buffer until it is retrieved by an XDREAD command with a starting LOGICAL BLOCK ADDRESS and TRANSFER LENGTH that match, or are a subset of, the LOGICAL BLOCK ADDRESS and REGENERATE LENGTH of this command.

4.0.17 REGENERATE (32) command

The REGENERATE (32) command (see table 54) requests that the target write to the buffer the XOR data generated from its own medium and the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data [and protection information, if any](#).

4.0.18 SYNCHRONIZE CACHE (10) command

The SYNCHRONIZE CACHE (10) command (see table 60) ensures that logical blocks in the cache memory, within the specified range, have their most recent data value [and protection information, if any](#), recorded on the physical medium. If a more recent data value for a logical block within the specified range exists in the cache memory than on the physical medium, then the logical block from the cache memory shall be written to the physical medium. Logical blocks may not be removed from the cache memory as a result of the synchronize cache operation. The synchronize cache function is also required implicitly by other SCSI functions as defined in other clauses of this standard.

4.0.19 SYNCHRONIZE CACHE (16) command

The SYNCHRONIZE CACHE (16) command (see table 61) ensures that logical blocks in the cache memory, within the specified range, have their most recent data value [and protection information, if any](#), recorded on the physical medium. If a more recent data value for a logical block within the specified range exists in the cache memory than on the physical medium, then the logical block from the cache memory shall be written to the physical medium. Logical blocks may not be removed from the cache memory as a result of the synchronize cache operation. The synchronize cache function is also required implicitly by other SCSI functions as defined in other clauses of this standard

4.0.20 VERIFY (10) command

The VERIFY (10) command (see table 8) requests that the device server verify the data written on the medium.

Table 8 — VERIFY (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (2Fh)							
1	VRPDATA	Reserved		DPO	Reserved	BLKVFY	BYTCHK	RELADR
2	(MSB)	LOGICAL BLOCK ADDRESS						
5								
6	Restricted for MMC-4	Reserved						
7	(MSB)	VERIFICATION LENGTH						
8								
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See 4.0.9 for a description of the cache control bit (DPO). See the LOCK UNLOCK CACHE (10) command (see 5.2.3) for a definition of the RELADR bit and the LOGICAL BLOCK ADDRESS field.

If the MODE SELECT command is implemented, and the Verify Error Recovery mode page is also implemented, then the current settings in that page specifies the verification criteria. If the Verify Error Recovery mode page is not implemented, then the verification criteria is vendor-specific.

If the byte check (BYTCHK) bit is zero, a medium verification shall be performed with no data comparison. If the BYTCHK bit is one, a byte-by-byte comparison of data written on the medium and the data transferred from the application client shall be performed. If the comparison is unsuccessful for any reason, the device server shall return CHECK CONDITION status and the sense key shall be set to MISCOMPARE with the appropriate additional sense code for the condition.

For direct access block devices, the blank verify (BLKVFY) bit shall be considered reserved. For optical and write-once block devices, the BLKVFY BIT is defined as follows. If the BLKVFY bit is zero, the device server shall not verify that the blocks are blank. If the BLKVFY bit is one, the device server shall verify that the blocks are blank. If the BYTCHK is one and the BLKVFY bit is one the device server shall return CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN CDB.

If the verify protection data (VRPDATA) is set to zero the protection information, if any, shall not be verified. If the VRPDATA is set to one the values in the other bits and fields in VERIFY command shall determine the method of verification of the protection information.

The VERIFICATION LENGTH field specifies the number of contiguous logical blocks of data or blanks that shall be verified. A VERIFICATION LENGTH of zero indicates that no logical blocks shall be verified. This condition shall not be considered as an error. Any other value indicates the number of logical blocks that shall be verified. The VERIFICATION LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.1.4.2).

4.0.21 VERIFY (12) command

The VERIFY (12) command (see table 9) requests that the device server verify the data on the medium.

Table 9 — VERIFY (12) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (AFh)							
1	VRPDATA	Reserved		DPO	Reserved	BLKVfy	BYCHK	RELADR
2	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
5								
6	(MSB)	VERIFICATION LENGTH						(LSB)
9								
10	Reserved							
11	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the VERIFY (10) command (see 4.0.20) for a description of the fields in this command.

4.0.22 VERIFY (16) command

The VERIFY (16) command (see table 10) requests that the device server verify the data written on the medium.

Table 10 — VERIFY (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (8Fh)							
1	VRPDATA	Reserved		DPO	Reserved	BLKVfy	BYTCHK	RELADR
2	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
9								
10	(MSB)	VERIFICATION LENGTH						(LSB)
13								
14	Reserved							
15	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the VERIFY (10) command (see 4.0.20) for a description of the fields in this command.

4.0.23 WRITE (10) command

The WRITE (10) command (see table 11) requests that the device server write the data transferred by the application client to the medium.

Table 11 — WRITE (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (2Ah)							
1	WRPROTECT	Reserved	DPO	FUA	EBP	Reserved	RELADR	
2	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
5		Reserved						
6	(MSB)	TRANSFER LENGTH						(LSB)
8		CONTROL						
9								

The [WRPROTECT](#) field is defined in table 12.

Table 12 — WRPROTECT field

Value	Description
00b	Preserve the contents of the protection fields (e.g., write to media, store in non-volatile memory, recalculate on read back). The device server shall determine if the data block is valid by checking the contents of the protection fields. If the device server determines there is an error the command shall be terminated with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to LOGICAL UNIT COMMUNICATION FAILURE. If the logical unit has not been formatted to accept protection information the device server shall terminate the command with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to CANNOT WRITE MEDIUM - UNKNOWN FORMAT.
01b	Preserve the contents of the protection fields (e.g., write to media, store in non-volatile memory, recalculate on read back). The device server shall not check the contents of the protection fields. If the logical unit has not been formatted to accept protection information the device server shall terminate the command with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to CANNOT WRITE MEDIUM - UNKNOWN FORMAT.
10b	The contents of the protection fields shall not be preserved. The device server shall determine if the data block is valid by checking the contents of the protection fields. If the device server determines there is an error the command shall be terminated with a CHECK CONDITION status and the sense data shall be set to MEDIUM ERROR with the additional sense code set to LOGICAL UNIT COMMUNICATION FAILURE. In the case where the logical unit has not been formatted with protection fields and the device server does not check the WRPROTECT field the device servers response to the write command is unpredictable.
11b	Reserved

4.0.24 WRITE (12) command

The WRITE (12) command (see table 13) requests that the device server write the data transferred from the application client to the medium.

Table 13 — WRITE (12) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (AAh)							
1	WRPROTECT	Reserved	DPO	FUA	Reserved		RELADR	
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	(MSB) _____							
9	TRANSFER LENGTH _____ (LSB)							
10	Restricted for MMC-4	Reserved						
11	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the WRITE (10) command (see 4.0.23) for a description of the fields in this command.

4.0.25 WRITE (16) command

The WRITE (16) command (see table 14) requests that the device server write the data transferred by the application client to the medium.

Table 14 — WRITE (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (8Ah)							
1	WRPROTECT		Reserved	DPO	FUA	Reserved		RELADR
2	(MSB) _____							
9	LOGICAL BLOCK ADDRESS _____ (LSB)							
10	(MSB) _____							
13	TRANSFER LENGTH _____ (LSB)							
14	Reserved							
15	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the WRITE (10) command (see 4.0.23) for a description of the fields in this command.

4.0.26 WRITE AND VERIFY (10) command

The WRITE AND VERIFY (10) command (see table 15) requests that the device server write the data transferred from the application client to the medium and then verify that the data [and protection information, if any](#), is correctly written. The data is only transferred once from the application client to the device server.

Table 15 — WRITE AND VERIFY (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (2Eh)							
1	WRPROTECT	Reserved	DPO	Reserved	EBP	BYTCHK	RELADR	
2	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
5		Reserved						
6	(MSB)	TRANSFER LENGTH						(LSB)
8		CONTROL						
9								

See 4.2.1.8 for reservation requirements for this command. See the LOCK UNLOCK CACHE (10) command (see 5.2.3) for a definition of the RELADR bit and the LOGICAL BLOCK ADDRESS field. See the WRITE (10) command (4.0.23) for a definition of the TRANSFER LENGTH field [and the WRPROTECT field](#). See 4.0.9 for a description of the cache control bit (DPO). See the WRITE (10) command (see 4.0.23) for a description of the EBP bit.

If the MODE SELECT command is implemented, and the Verify Error Recovery mode page is also implemented, then the current settings in that mode page (along with the AWRE bit from the Read-Write Error Recovery mode page) specify the verification error criteria. If these mode pages are not implemented, then the verification criteria is vendor-specific.

A byte check (BYTCHK) bit of zero requests a medium verification to be performed with no data comparison. A BYTCHK bit of one requests a byte-by-byte comparison of data written on the medium and the data transferred from the application client. If the comparison is unsuccessful for any reason, the device server shall return CHECK CONDITION status with the sense key set to MISCOMPARE with the appropriate additional sense code for the condition.

Editor's Note 1: Add in a new ASCQ of 0Ch 0Eh titled PROTECTION INFORMATION WRITE ERROR.

NOTE 2 - The WRITE AND VERIFY command specifically states that the data are not to be transferred twice (i.e., once for the write pass, and once for the verify pass) when performing a byte compare. If there is a need for two transfers to occur (e.g., to ensure the integrity of the path to the media), then the application client should issue a WRITE command with a LINK bit of one followed by a VERIFY command with a BYTCHK bit of one, transferring the same data on each command.

4.0.27 WRITE AND VERIFY (12) command

The WRITE AND VERIFY (12) command (see table 16) requests that the device server write the data transferred from the application client to the medium and then verify that the data [and protection information, if any](#), is correctly written.

Table 16 — WRITE AND VERIFY(12) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (AEh)							
1	WRPROTECT	Reserved		DPO	Reserved	EBP	BYCHK	RELADR
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	(MSB) _____							
9	TRANSFER LENGTH _____ (LSB)							
10	Reserved							
11	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the WRITE AND VERIFY (10) command (see 4.0.26) for a description of the bits in this command.

4.0.28 WRITE AND VERIFY (16) command

The WRITE AND VERIFY (16) command (see table 17) requests that the device server write the data transferred from the application client to the medium and then verify that the data [and protection information, if any](#), is correctly written. The data is only transferred once from the application client to the device server.

Table 17 — WRITE AND VERIFY (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (8Eh)							
1	WRPROTECT		Reserved	DPO	Reserved	EBP	BYCHK	RELADR
2	(MSB) _____							
9	LOGICAL BLOCK ADDRESS _____ (LSB)							
10	(MSB) _____							
13	TRANSFER LENGTH _____ (LSB)							
14	Reserved							
15	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the WRITE AND VERIFY (10) command (see 4.0.26) for a description of the fields in this command.

4.0.29 WRITE LONG command

The WRITE LONG command (see table 72) requests that the device server write the data transferred by the application client to the medium. The data passed during the WRITE LONG command is implementation specific, but shall include the data bytes, [any protection information](#), and the ECC bytes.

4.0.30 WRITE SAME (10) command

The WRITE SAME (10) command (see table 18) requests that the device server write the single block of data transferred by the application client to the medium multiple times to consecutive multiple logical blocks. [If the medium is formatted with protection information the value in the LOGICAL BLOCK ADDRESS field \(see 4.0.2\) shall be placed into the LOGICAL ADDRESS TAG field \(see 4.0.2\) of the first logical block written to the medium. Into each of the following logical blocks the LOGICAL ADDRESS TAG field shall be placed the logical block address received in the WRITE SAME command incremented by one.](#)

Table 18 — WRITE SAME (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (41h)							
1	WRPROTECT		Reserved			PBDATA	LBDATA	RELADR
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	Reserved							
7	(MSB) _____							
8	NUMBER OF BLOCKS _____ (LSB)							
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the LOCK UNLOCK CACHE (10) command (see 5.2.3) for a definition of the RELADR bit and the LOGICAL BLOCK ADDRESS field. [See the WRITE \(10\) command \(4.0.23\) for a definition of the WRPROTECT field.](#)

4.0.31 WRITE SAME (16) command

The WRITE SAME (16) command (see table 19) requests that the device server write the single block of data transferred by the application client to the medium multiple times to consecutive multiple logical blocks. [If the medium is formatted with protection information the value in the LOGICAL BLOCK ADDRESS field \(see 4.0.2\) shall be placed into the LOGICAL ADDRESS TAG field \(see 4.0.2\) of the first logical block written to the medium. Into each of the following logical blocks the LOGICAL ADDRESS TAG field shall be placed the logical block address received in the WRITE SAME command incremented by one.](#)

Table 19 — WRITE SAME (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (93h)							
1	WRPROTECT		Reserved			PBDATA	LBDATA	RELADR
2	(MSB) _____							
9	LOGICAL BLOCK ADDRESS _____ (LSB)							
10	(MSB) _____							
13	NUMBER OF BLOCKS _____ (LSB)							
14	Reserved							
15	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the WRITE SAME (10) command (see 4.0.30) for a description of the fields in this command.

4.0.32 XDREAD (10) command

The XDREAD (10) command (see table 20) requests that the target transfer to the initiator the XOR data generated by an XDWRITE or REGENERATE command.

Table 20 — XDREAD (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (52h)							
1	XORPDATA	Reserved						
2	(MSB)	LOGICAL BLOCK ADDRESS						
5								
6	Reserved							
7	(MSB)	TRANSFER LENGTH						
8								
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command.

[If the XOR protection data \(XORPDATA\) is set to zero the protection information, if any, shall not be verified or transmitted. If the XORPDATA is set to one the protection information shall be transmitted but shall not be verified.](#)

The XOR data transferred is identified by the LOGICAL BLOCK ADDRESS and TRANSFER LENGTH. The LOGICAL BLOCK ADDRESS and TRANSFER LENGTH shall be the same as, or a subset of, those specified in a prior XDWRITE or REGENERATE command. If a match is not found the command is terminated with a CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN CDB. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.1.4.2).

4.0.33 XDREAD (32) command

The XDREAD (32) command (see table 21) requests that the target transfer to the initiator the XOR data generated by an XDWRITE or REGENERATE command.

Table 21 — XDREAD (32) command

Byte\Bit	7	6	5	4	3	2	1	0							
0	OPERATION CODE (7Fh)														
1	CONTROL														
2	Reserved														
6															
7	ADDITIONAL CDB LENGTH (18h)														
8	(MSB)	SERVICE ACTION (0003h)						(LSB)							
9															
10	XORPDATA	Reserved													
11	Reserved														
12	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)							
19															
20	Reserved														
27															
28	(MSB)	TRANSFER LENGTH						(LSB)							
31															

See 4.2.1.8 for reservation requirements for this command. See the XDREAD (10) command (see 4.0.32) and SPC-3 for a description of the fields in this command.

4.0.34 XDWRITE (10) command

The XDWRITE (10) command (see table 22) requests that the target XOR the data transferred [including the protection information, if any](#), with the data on the medium. The resulting XOR data is stored by the target until it is retrieved by an XDREAD (10) command.

Table 22 — XDWRITE (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (50h)							
1	WRPROTECT		Reserved	DPO	FUA	DISABLE WRITE	Reserved	
2	(MSB) _____ LOGICAL BLOCK ADDRESS _____ (LSB)							
5								
6	Reserved							
7	(MSB) _____ TRANSFER LENGTH _____ (LSB)							
8								
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the READ (10) command (see 4.0.9) for a definition of the cache control bits (DPO and FUA). [See the WRITE \(10\) command \(4.0.23\) for a definition of the WRPROTECT field.](#)

4.0.35 XDWRITE (32) command

The XDWRITE (32) command (see table 23) requests that the target XOR the data transferred [including the protection information, if any](#), with the data on the medium. The resulting XOR data is stored by the target until it is retrieved by an XDREAD (32) command.

Table 23 — XDWRITE (32) command

Byte\Bit	7	6	5	4	3	2	1	0					
0	OPERATION CODE (7Fh)												
1	CONTROL												
2	Reserved												
6													
7	ADDITIONAL CDB LENGTH (18h)												
8	(MSB)	SERVICE ACTION (0004h)						(LSB)					
9													
10	WRPROTECT	Reserved	DPO	FUA	DISABLE WRITE	Reserved							
11	Reserved												
12	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)					
19													
20	Reserved												
27													
28	(MSB)	TRANSFER LENGTH						(LSB)					
31													

See 4.2.1.8 for reservation requirements for this command. See the XDWRITE (10) command (see 4.0.34) and SPC-3 for a description of the fields in this command.

4.0.36 XDWRITEREAD (10) command

The XDWRITEREAD (10) command (see table 24) requests that the target XOR the data transferred (data-out) [including the protection information, if any](#), with the data on the medium and return the resulting XOR data (data-in). This is the equivalent to an XDWRITE (10) followed by an XDREAD (10) with the same

logical block address and transfer length. This command is only available on transport protocols supporting bidirectional commands.

Table 24 — XDWRITEREAD (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (53h)							
1	WRPROTECT		XORPDATA	DPO	FUA	DISABLE WRITE	Reserved	
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	Reserved							
7	(MSB) _____							
8	TRANSFER LENGTH _____ (LSB)							
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See the XWRITE (10) command (see 4.0.34) and XDREAD (10) command (see 4.0.32) for a description of the fields in this command.

4.0.37 XDWRITEREAD (32) command

The XDWRITEREAD (32) command (see table 25) requests that the target XOR the data transferred (data-out) [including the protection information, if any](#), with the data on the medium and return the resulting XOR data (data-in). This is the equivalent to an XDWRITE (32) followed by an XDREAD (32) with the same

logical block address and transfer length. This command is only available on transport protocols supporting bidirectional commands.

Table 25 — XDWRITEREAD (32) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (7Fh)							
1	CONTROL							
2	Reserved							
6								
7	ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (0007h)						(LSB)
9								
10	WRPROTECT	XORPDATA	DPO	FUA	DISABLE WRITE	Reserved		
11	Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
19								
20	Reserved							
27								
28	(MSB)	TRANSFER LENGTH						(LSB)
31								

See 4.2.1.8 for reservation requirements for this command. See the XDWRITEREAD (10) command (see 4.0.36) and SPC-3 for a description of the fields in this command.

4.0.38 XDWRITE EXTENDED (16) command

The XDWRITE EXTENDED (16) command (see table 26) requests that the target XOR the data transferred including the protection information, if any, with the data on the medium. The resulting XOR data including the protection information, if any, may subsequently be sent to a secondary device using an XPWRITE (10) or XPWRITE (32) command. The target, acting as a temporary initiator, issues XPWRITE commands to retrieve the specified data. XPWRITE (16) should be used for access to SCSI devices supporting less than 2

Terabytes, and XPWRITE (32) should be used for accesses to SCSI devices supporting greater than or equal to 2 Terabytes.

Table 26 — XDWRITE EXTENDED (16) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (80h)							
1	WRPROTECT		Reserved	DPO	FUA	DISABLE WRITE	PORT CONTROL	
2	(MSB) LOGICAL BLOCK ADDRESS (LSB)							
5								
6	(MSB) SECONDARY LOGICAL BLOCK ADDRESS (LSB)							
9								
10	(MSB) TRANSFER LENGTH (LSB)							
13								
14	SECONDARY ADDRESS							
15	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See 4.0.9 for a definition of the DPO and FUA bits. [See the WRITE \(10\) command \(4.0.23\) for a definition of the WRPROTECT field.](#)

4.0.39 XDWRITE EXTENDED (32) command

The XDWRITE EXTENDED (32) command (see table 27) requests that the target XOR the data transferred including the protection information, if any, with the data on the medium. The resulting XOR data may subsequently be sent to a secondary device using an XPWRITE (32) command.

Table 27 — XDWRITE EXTENDED (32) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (7Fh)							
1	CONTROL							
2	Reserved							
6								
7	ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (0005h)						(LSB)
9								
10	WRPROTECT	Reserved	DPO	FUA	DISABLE WRITE	PORT CONTROL		
11	SECONDARY ADDRESS							
12	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)
19								
20	(MSB)	SECONDARY LOGICAL BLOCK ADDRESS						(LSB)
27								
28	(MSB)	TRANSFER LENGTH						(LSB)
31								

See 4.2.1.8 for reservation requirements for this command. See the XWRITE EXTENDED (16) command (see 4.0.38) and SPC-3 for a description of the fields in this command.

4.0.40 XDWRITE EXTENDED (64) command

The XDWRITE EXTENDED (64) command (see table 28) requests that the target XOR the data transferred including the protection information, if any, with the data on the medium. The resulting XOR data may subsequently be sent to a secondary device using an XPWRITE (32) command.

Table 28 — XDWRITE EXTENDED (64) command

Byte\Bit	7	6	5	4	3	2	1	0	
0	OPERATION CODE (7Fh)								
1	CONTROL								
2	Reserved								
6									
7	ADDITIONAL CDB LENGTH (18h)								
8	(MSB)	SERVICE ACTION (0005h)						(LSB)	
9									
10	WRPROTECT	Reserved	DPO	FUA	DISABLE WRITE	PORT CONTROL			
12	SECONDARY ADDRESS DESCRIPTOR								
43									
44	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)	
51									
52	(MSB)	SECONDARY LOGICAL BLOCK ADDRESS						(LSB)	
59									
60	(MSB)	TRANSFER LENGTH						(LSB)	
63									

See 4.2.1.8 for reservation requirements for this command.

The SECONDARY ADDRESS DESCRIPTOR field contains the logical unit identifier of the logical unit that will receive the XOR data transfer. The format of this field conforms to one of the target descriptor formats of the EXTENDED COPY command as specified in SPC-3.

See the XWRITE EXTENDED (16) command (see 4.0.38) and SPC-3 for a description of the other fields in this command.

4.0.41 XPWRITE (10) command

The XPWRITE (10) command (see table 29) requests that the target XOR the data transferred [including the protection information, if any](#), with the data on the medium and then write the XOR data to the medium.

Table 29 — XPWRITE (10) command

Byte\Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (51h)							
1	WRPROTECT		Reserved	DPO	FUA	Reserved		
2	(MSB) _____							
5	LOGICAL BLOCK ADDRESS _____ (LSB)							
6	Reserved							
7	(MSB) _____							
8	TRANSFER LENGTH _____ (LSB)							
9	CONTROL							

See 4.2.1.8 for reservation requirements for this command. See 4.0.9 for a definition of the DPO and FUA bits. [See the WRITE \(10\) command \(4.0.23\) for a definition of the WRPROTECT field.](#)

The LOGICAL BLOCK ADDRESS field specifies the starting logical block address where the target shall read data from its medium. It also specifies the starting logical block address where the XOR result data shall be written to the medium.

The TRANSFER LENGTH field specifies the number of blocks that shall be read from the medium. It also specifies the number of blocks that shall be written to the medium. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.1.4.2).

4.0.42 XPWRITE (32) command

The XPWRITE (32) command (see table 30) requests that the target XOR the data transferred [including the protection information, if any](#), with the data on the medium and then write the XOR data to the medium.

Table 30 — XPWRITE (32) command

Byte\Bit	7	6	5	4	3	2	1	0							
0	OPERATION CODE (7Fh)														
1	CONTROL														
2	Reserved														
6															
7	ADDITIONAL CDB LENGTH (18h)														
8	(MSB)	SERVICE ACTION (0006h)						(LSB)							
9															
10	WRPROTECT	Reserved	DPO	FUA	Reserved										
44	(MSB)	LOGICAL BLOCK ADDRESS						(LSB)							
51															
52	Reserved														
59															
60	(MSB)	TRANSFER LENGTH						(LSB)							
63															

See 4.2.1.8 for reservation requirements for this command. See the XPWRITE (10) command (see 4.0.41) and SPC-3 for a description of the fields in this command.

SPC-3 additions

4.0.43 Standard INQUIRY data

The standard INQUIRY data (see table 31) shall contain at least 36 bytes.

Table 31 — Standard INQUIRY data format

Bit Byte	7	6	5	4	3	2	1	0
0	PERIPHERAL QUALIFIER			PERIPHERAL DEVICE TYPE				
1	RMB	Reserved						
2	VERSION							
3	Obsolete	Obsolete	NORMACA	HiSUP	RESPONSE DATA FORMAT			
4	ADDITIONAL LENGTH (n-4)							
5	SCCS	ACC	ALUA		3PC	Reserved	PROTECT	
6	BQUE	ENC SERV	VS	MULTIP	MCHNGR	Obsolete	Obsolete	ADDR16 ^a
7	RELADR	Obsolete	WBUS16 ^a	SYNC ^a	LINKED	Obsolete	CMDQUE	VS
8	(MSB) _____							
15	VENDOR IDENTIFICATION _____ (LSB)							
16	(MSB) _____							
31	PRODUCT IDENTIFICATION _____ (LSB)							
32	(MSB) _____							
35	PRODUCT REVISION LEVEL _____ (LSB)							
36	_____							
55	Vendor specific _____							
56	Reserved				CLOCKING ^a		QAS ^a	IUS ^a
57	Reserved							
58	(MSB) _____							
59	VERSION DESCRIPTOR 1 _____ (LSB)							
	⋮							
72	(MSB) _____							
73	VERSION DESCRIPTOR 8 _____ (LSB)							
74	_____							
95	Reserved _____							
	Vendor specific parameters							
96	_____							
n	Vendor specific _____							

^a The meanings of these fields are specific to SPI-5 (see 6.4.3). For SCSI protocols other than the SCSI Parallel Interface, these fields are reserved.

[The PROTECT field is defined in table 32](#)

Table 32 — Peripheral qualifier

Qualifier	Description
00b	This logical unit does not support the protection fields (see SBC-2).
01b	This logical unit supports the protection fields but the medium has not been formatted to include protection fields.
10b	Reserved
11b	The medium for this logical unit has been formatted to include protection fields with each logical block.

4.1 EXTENDED COPY command

4.1.1 EXTENDED COPY command introduction

4.1.2 Errors detected before starting processing of the segment descriptors

4.1.3 Errors detected during processing of segment descriptors

4.1.4 Abort task management functions

4.1.5 Descriptor type codes

4.1.6 Target descriptors

4.1.6.1 Target descriptors introduction

4.1.6.2 Identification descriptor target descriptor format

4.1.6.3 Alias target descriptor format

4.1.6.4 Device type specific target descriptor parameters for block device types

The format for the device type specific target descriptor parameters for block device types (device type code values 00h, 04h, 05h, 07h, and 0Eh) is shown in table 33.

Table 33 — Device type specific target descriptor parameters for block device types

Bit Byte	7	6	5	4	3	2	1	0
28	PROTECT_ON	Reserved				PAD	Reserved	
29	(MSB)							
30	DISK BLOCK LENGTH							
31	(LSB)							

The PAD bit is used in conjunction with the CAT bit (see 4.1.7.1) in the segment descriptor to determine what action should be taken when a segment of the copy does not fit exactly into an integer number of destination blocks.

[The protection information enabled \(PROTECT_ON\) bit is set to zero indicates the block device has not been formatted to include protection fields. The PROTECT_ON bit set to one indicates the block device has been formatted to include protection fields with each logical block.](#)

The DISK BLOCK LENGTH field contains the number of bytes in a disk block for the logical device being addressed.

The copy manager may read ahead from sources of block device type. That is, the copy manager may perform read operations from a source disk at any time and in any order during processing of an EXTENDED COPY command, provided that the relative order of writes and reads on the same blocks within the same target descriptor does not differ from their order in the segment descriptor list.

4.1.6.5 Device type specific target descriptor parameters for sequential-access device types

The format for the device type specific target descriptor parameters for the sequential-access device type (device type code value 01h) is shown in table 34.

Table 34 — Device type specific target descriptor parameters for sequential-access device types

Bit Byte	7	6	5	4	3	2	1	0
28	PROTECT_ON	Reserved				PAD	Reserved	FIXED
29	(MSB)							
30		STREAM BLOCK LENGTH						
31		(LSB)						

The contents of the FIXED bit and STREAM BLOCK LENGTH field are combined with the STREAM DEVICE TRANSFER LENGTH FIELD in the segment descriptor to determine the length of the stream read or write operation as specified in table 35.

Table 35 — Stream device transfer lengths

FIXED bit	STREAM BLOCK LENGTH field	Description
0	0	Use variable length reads or writes. The number of bytes for each read or write is specified by the STREAM DEVICE TRANSFER LENGTH field in the segment descriptor.
0	not 0	The command shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.
1	0	The command shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.
1	not 0	Use fixed record length reads or writes. The number of bytes for each read or write shall be the product of the STREAM BLOCK LENGTH field and the STREAM DEVICE TRANSFER LENGTH field in the segment descriptor.

The PAD bit is used in conjunction with the CAT bit (see 4.1.7.1) in the segment descriptor to determine what action should be taken when a segment of the copy does not fit exactly into an integer number of destination blocks.

[The protection information enabled \(PROTECT_ON\) bit set to zero indicates the block device has not been formatted to include protection fields. The PROTECT_ON bit set to one indicates the block device has been formatted to include protection fields with each logical block. If the PROTECT_ON bit is set to one and the FIXED](#)

[bit is set to zero, or the FIXED bit is set to one and the STREAM BLOCK LENGTH field is set to zero then the command shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN PARAMETER LIST.](#)

All read commands issued to sequential-access type devices shall have the SILI bit equal to zero.

NOTE 3 - It is anticipated that future versions of this standard may use bit 1 of byte 28 in the device type specific target descriptor parameters for stream device types to indicate the value of the SILI bit for read commands, after T10 establishes how the copy manager is to process tape reads of unknown block length without error.

The copy manager shall not read ahead from sources of stream device type. That is, the read operations required by a segment descriptor for which the source is a stream device shall not be started until all write operations for previous segment descriptors have completed.

4.1.6.6 Device type specific target descriptor parameters for processor device types

4.1.7 Segment Descriptors

4.1.7.1 Segment descriptors introduction

All segment descriptors begin with the eight byte header shown in table 36.

Table 36 — Segment descriptor header

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (00h-3Fh)							
1	CK_PROTECT	Reserved					DC	CAT
2	(MSB)	DESCRIPTOR LENGTH						(LSB)
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						(LSB)
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						(LSB)
7								

The descriptor type code field is described in 4.1.5. Support for each segment descriptor format is optional. If copy manager receives an unsupported descriptor type code in a segment descriptor, the command shall be terminated with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to UNSUPPORTED SEGMENT DESCRIPTOR TYPE CODE.

[If the PROTECT_ON bit is set to zero \(see 4.1.6.4\) the ck_protect bit shall be ignored.](#)

[The check protection \(CK_PROTECT\) bit set to zero indicates the copy manager shall issue commands that prevent the contents of the protection fields being checked. If the CK_PROTECT bit is set to one copy manager shall issue commands that cause the contents of the protection fields to be checked.](#)

4.1.7.2 Block device to stream device operations

The segment descriptor format shown in table 37 is used by the copy operations that move data from a block device to a stream device or vice versa.

Table 37 — Block device to or from stream device segment descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (00h, 01h, 0Bh, or 0Ch)							
1	CK_PROTECT	Reserved						CAT
2	(MSB)	DESCRIPTOR LENGTH (0014h)						
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						
7								
8	Reserved							
9	(MSB)	STREAM DEVICE TRANSFER LENGTH						
10								
11								
12	Reserved							
13	Reserved							
14	(MSB)	BLOCK DEVICE NUMBER OF BLOCKS						
15								
16	(MSB)	BLOCK DEVICE LOGICAL BLOCK ADDRESS						
23								

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.3 Stream device to block device operations

4.1.7.4 Block device to block device operations

The segment descriptor format shown in table 38 is used by the copy operations that move data from a block device to a block device.

Table 38 — Block device to block device segment descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (02h or 0Dh)							
1	CK_PROTECT	Reserved					DC	CAT
2	(MSB)	DESCRIPTOR LENGTH (0018h)						
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						
7								
8	Reserved							
9	Reserved							
10	(MSB)	BLOCK DEVICE NUMBER OF BLOCKS						
11								
12	(MSB)	SOURCE BLOCK DEVICE LOGICAL BLOCK ADDRESS						
19								
20	(MSB)	DESTINATION BLOCK DEVICE LOGICAL BLOCK ADDRESS						
27								

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.5 Stream device to stream device operations

The segment descriptor format shown in table 39 is used by the copy operations that move data from a stream device to a stream device.

Table 39 — Stream device to stream device segment descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (03h or 0Eh)							
1	CK_PROTECT	Reserved						CAT
2	(MSB)	DESCRIPTOR LENGTH (0010h)						
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						
7								
8	Reserved							
9	(MSB)	SOURCE STREAM DEVICE TRANSFER LENGTH						
10								
11		Reserved						
12								
13	(MSB)	DESTINATION STREAM DEVICE TRANSFER LENGTH						
14								
15		BYTE COUNT						
16	(MSB)							
19		(LSB)						

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.6 Inline data to stream device operation

4.1.7.7 Embedded data to stream device operation

4.1.7.8 Stream device to discard operation

The segment descriptor format shown in table 40 instructs the copy manager to read data from a stream device and not copy it to any destination device.

Table 40 — Stream device to discard segment descriptor

Bit Byte	7	6	5	4	3	2	1	0					
0	DESCRIPTOR TYPE CODE (06h or 0Fh)												
1	CK_PROTECT	Reserved						CAT					
2	(MSB)	DESCRIPTOR LENGTH (000Ch)											
3									(LSB)				
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX											
5									(LSB)				
6	Reserved												
7	Reserved												
8	Reserved												
9	(MSB)	STREAM DEVICE TRANSFER LENGTH											
10													
11									(LSB)				
12	(MSB)	NUMBER OF BYTES											
15									(LSB)				

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.9 Verify device operation

4.1.7.10 Block device with offset to stream device operation

The segment descriptor format shown in table 41 is used to instruct the copy manager to move data from a block device with a byte offset to a stream device or vice versa.

Table 41 — Block device with offset to or from stream device segment descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (08h or 09h)							
1	CK_PROTECT	Reserved						CAT
2	(MSB)	DESCRIPTOR LENGTH (0018h)						
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						
7								
8	Reserved							
9	(MSB)	STREAM DEVICE TRANSFER LENGTH						
10								
11		NUMBER OF BYTES						
12	(MSB)							
15		BLOCK DEVICE LOGICAL BLOCK ADDRESS						
16	(MSB)							
23		Reserved						
24								
25		Reserved						
26	(MSB)							
27		BLOCK DEVICE BYTE OFFSET						(LSB)

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.11 Stream device to block device with offset operation

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.12 Block device with offset to block device with offset operation

The segment descriptor format shown in table 42 instructs the copy manager to move data from a block device with a byte offset to a block device with a byte offset.

Table 42 — Block device with offset to block device with offset segment descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE CODE (0Ah)							
1	CK_PROTECT	Reserved						CAT
2	(MSB)	DESCRIPTOR LENGTH (001Ch)						
3								
4	(MSB)	SOURCE TARGET DESCRIPTOR INDEX						
5								
6	(MSB)	DESTINATION TARGET DESCRIPTOR INDEX						
7								
8	(MSB)	NUMBER OF BYTES						
11								
12	(MSB)	SOURCE BLOCK DEVICE LOGICAL BLOCK ADDRESS						
19								
20	(MSB)	DESTINATION BLOCK DEVICE LOGICAL BLOCK ADDRESS						
27								
28	(MSB)	SOURCE BLOCK DEVICE BYTE OFFSET						
29								
30	(MSB)	DESTINATION BLOCK DEVICE BYTE OFFSET						
31								

...

[The CK_PROTECT bit is described in 4.1.7.1.](#)

4.1.7.13 Write filemarks operation**4.1.7.14 Space operation****4.1.7.15 Locate operation****4.1.7.16 Tape device image copy operation****4.1.7.17 Register key operation****4.1.8 Control mode page**

The Control mode page (see table 43) provides controls over several SCSI features that are applicable to all device types such as tagged queuing and error logging.

Table 43 — Control mode page

Bit Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (0Ah)					
1	PAGE LENGTH (0Ah)							
2	TST			Reserved	WR_CRC	D_SENSE	GLTSD	RLEC
3	QUEUE ALGORITHM MODIFIER				Reserved	QERR		DQUE
4	TAS	RAC	UA_INTLCK_CTRL		SWP	Obsolete		
5	Reserved					AUTOLOAD MODE		
6	Obsolete							
7								
8	(MSB)	BUSY TIMEOUT PERIOD						
9								(LSB)
10	(MSB)	EXTENDED SELF-TEST COMPLETION TIME						
11								(LSB)

A write CRC ([WR_CRC](#)) bit set to zero indicates the contents of the protected data blocks CRC field (see SBC-2) shall be preserved and may be written to media. A [WE_CRC](#) bit set to one indicates the contents of the protected data blocks CRC field shall be preserved by writing it to media. If the logical unit has not been formatted to include protection fields the [wr_crc](#) bit shall be ignored.