memorandum

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

T10/07-409r1

To	From	Subject	Date
INCITS T10 Committee	Michael Banther, HP	SSC-3 Volume Model	18 September 2007

Revision History

Revision 0 – Initial document.

Revision 1 – Modified the definition of a volume.

Reference documents

SCSI Stream Commands – 3 (SSC-3), Project 1611-D, Rev 03d, 9 July 2007.

Cleaning Model, T10/07-219r1, 11 July 2007.

Background

The SSC-3 model clause includes a description of a generic recordable volume. It does not include a description of a non-recordable volume (e.g. a cleaning cartridge) or of a special-purpose recordable volume (e.g. a microcode upgrade volume). The definition for 'volume' in SSC-3 limits it to a cartridge containing a recordable medium.

The lack of description for cleaning and microcode upgrade volumes leads to difficulties specifying the standard behaviour of an SSC-3 device server when interacting with one of these volumes. In particular, interoperability problems have occurred due to a different understanding of the concept of 'mounted' for cleaning and microcode upgrade volumes.

This proposal contains changes to the *definition* clauses (3.1), *physical elements* model clause (4.2.2), and document-wide usage of various related terms needed to include the cleaning volume type. It also contains changes to use of the terms:

a) 'volume' (149 instances) and 'volumes' (6 instances) to differentiate text that applies only to a volume containing a recordable medium from text that applies to any volume, whether containing a recordable medium or a cleaning medium;

b) 'cleaning' (20 instances), 'cleaning cycle' (1 instance), and 'head cleaning' (2 instances) which I have changed to 'cleaning operation'; and

c) 'cleaning media' (4 instances) and 'cleaning tape' (5 instances) which I have changed to 'cleaning volume'.

Due to the limited period of time before last technical input for SSC-3, this proposal does not include changes to define a microcode upgrade volume type. For the same reason, it does not include changes to the various instances of 'media' (125 instances) or 'medium' (706 instances) in the present draft standard that may refer to a volume containing a recordable medium.

Changes to the SSC-3 draft standard

3.1.5 beginning-of-medium (BOM): The extreme position along the medium in the direction away from the supply reel ~~that is accessible by the device~~ where an operation between the read/write mechanism and the medium occurs. This position may not coincide with a beginning-of-partition position.

3.1.X cleaning volume: A volume (see 3.1.81) containing a cleaning medium (see 4.2.2.3).

3.1.X+1 data volume: A volume (see 3.1.81) containing a recording medium (see 4.2.2.2).

3.1.19 end-of-medium (EOM): The extreme position along the medium in the direction away from the take-up reel ~~that is accessible by the device~~ where an operation between the read/write mechanism and the medium occurs. This position may not coincide with a end-of-partition position.

3.1.27 format label: A vendor-specific series of logical objects that contain information used to identify the data volume or data set.

3.1.52 partition: The entire usable region for recording and reading in a data volume or in a portion of a data volume, defined in a vendor-specific or format-specific manner (see 4.2.5).

3.1.X+2 recorded volume: A data volume (see 3.1.X+1) upon which user data has been recorded.

3.1.73 tape: The medium ~~on which data is recorded~~ that interacts with the read/write mechanism. The medium is normally a long thin medium that is spooled onto one or two reels, possibly within a cassette or cartridge.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

3.1.75 thread: A part of the loading process in which the ~~recording~~ medium is being engaged for positioning on a suitable transport mechanism (e.g., spooled on to a take up reel, wrapped around the surface of a helical scan drum). After threading is complete the tape device may beginning positioning the medium to an initial position.

3.1.81 unthread: A part of the unloading process in which the ~~recording~~ medium is being disengaged from the suitable transport mechanism (e.g., de-spooled from a take up reel, unwrapped from around the surface of a helical scan drum).

3.1.83 vendor-specific control meta-data: Vendor-specific information stored on the data volume outside the user data area(s) that is used to control or specify how the data volume is being used by application clients (e.g., directory information, partition information, EOD locations, copies of data stored in a vendor-specific manner, volume serial number information, number of logical blocks on the media).

3.1.84 volume: A ~~recording~~ removable ~~medium together with its~~ physical carrier containing a medium.

4.2.2 ~~Physical elements~~ Removable medium

4.2.2.1 Removable medium introduction

The ~~recording~~ medium for tape devices consists of various widths and lengths of a flexible substrate. All tape devices use a recording medium for storage of data and some tape devices also use a cleaning medium for cleaning the read/write mechanism. A recording medium has the substrate coated with a semi-permanent magnetic material. A cleaning medium uses an abrasive substrate. The ~~recording~~ medium may be spooled onto single reels or encapsulated into cartridges containing both a supply reel and a take-up reel. Several American National Standards exist covering the construction of reels and cartridges for interchange as well as recording techniques for many of the format or density combinations.

For a sequential-access device, a ~~recording~~ medium exists between two reels, the supply reel and take-up reel. The read/write mechanism ~~may~~ only ~~access~~ interacts with the medium between the reels. As the medium is taken out of one reel, it passes by the read/write mechanism and into the other reel. Transferring data as a stream is most efficient, since the recording medium ~~may~~ traverses the read/write mechanism producing a flow of data. To position to a given point requires moving the medium until the appropriate position is found.

The ~~recording~~ medium has two physical attributes called beginning-of-medium (BOM) and end-of-medium (EOM). Beginning-of-medium is at the end of the medium that is attached to the take-up reel. End-of-medium is at the end of the medium that is attached to the supply reel. In some cases, the medium is permanently affixed to one or both of the reel hubs. ~~Beginning or end~~ For a recording medium, the beginning of medium and the end of medium ~~is not required to be related~~ may not have a specific relationship to the beginning or end of any partition.

A volume is composed of the ~~recording~~ medium and its physical carrier (e.g., reel, cartridge, cassette). Volumes have an attribute of being mounted or de-mounted on a suitable transport mechanism.

Mounted is the state of a volume when the device ~~is physically capable of processing~~ moves the medium past the read/write mechanism to process commands, ~~that cause the medium to be moved~~ to move to a position, or in the case of a cleaning volume to clean the read/write mechanism. A volume is de-mounted when it is being loaded, threaded, unloaded, unthreaded, or when not attached to the device.

Ready is the state of the logical unit when the device server processes medium access and non-medium access commands ~~may be processed~~. The logical unit is not ready when no volume is mounted or, from the SCSI initiator device perspective, whenever any medium access command reports CHECK CONDITION status and a NOT READY sense key. The logical unit is not ready during the transition from mounted to not mounted, or not mounted to mounted. Devices may have a physical control that places the device in a not ready state even when a volume is mounted.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

4.2.2.2 Recording medium model

As shown in figure 3, a portion of the physical length of recording medium is not usable for recording data. For most data volumes, a length of the recording medium is reserved between the take-up reel and the beginning-of-medium, and between the end-of-medium position and the supply reel. This is done to provide a sufficient tape wrap onto the reel hub and to ensure that recording starts in an undamaged section of the medium.

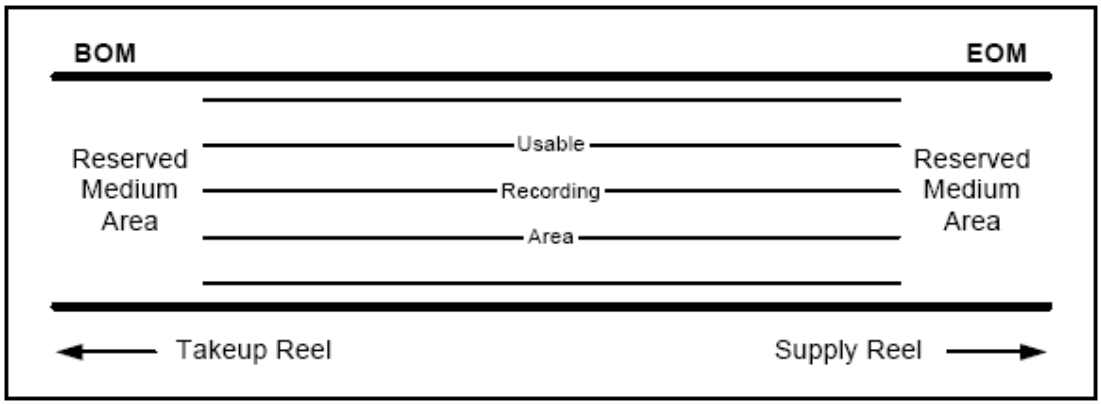

Figure 3 — Typical ~~volume~~ recording medium layout

The position on the recording medium where one write component records a pattern of ~~recorded~~ signals ~~may be written by one write component~~ is called a track (see figure 4). A device ~~may~~ writes or reads from one or more tracks at a time, depending on the format.

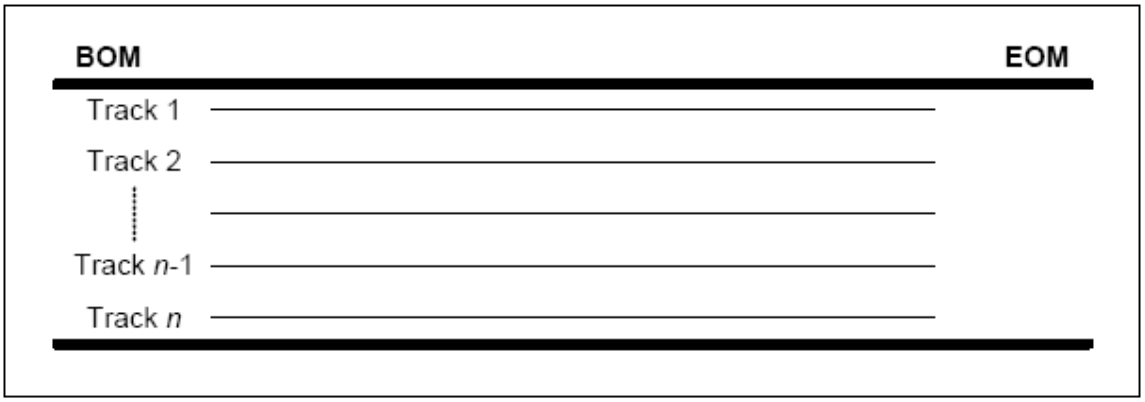

Figure 4 — Typical recording medium track layout

On a new data volume, recording of one or more tracks begins after mounting the data volume and moves from beginning-of-medium toward end-of-medium. The number of tracks written at one time is called a track group (TrkGrp). The use of ~~T~~track groups ~~may be used by~~ is independent of any recording format. For recorded data volumes, reading in the forward direction follows the same course of tracks when writing.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

In serpentine recording, not all tracks are recorded at the same time. At the end-of-medium or beginning-of-medium, the device reverses direction and begins recording the next track group. The process of reversing direction and recording the next track group ~~may be repeated~~ repeats until all track groups are recorded. For serpentine devices that record only one track at a time, each physical track represents one track group (see figure 5).

Figure 5 — Serpentine recording example

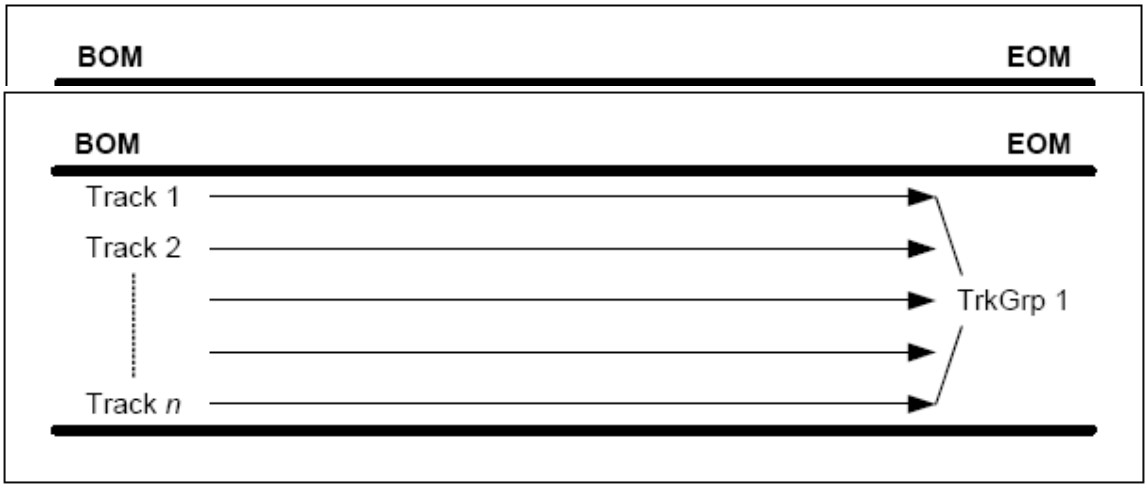

Some multi-track devices have only one track group, using a parallel storage format that supports the simultaneous recording of all available tracks (see figure 6).

Figure 6 — Parallel recording example

The serpentine and parallel recording formats shown in the previous examples define tracks as longitudinal patterns of recorded information. One other storage format used by some devices records tracks diagonally across the recording medium. Recording of ~~O~~one or more tracks ~~may be recorded~~ occurs at the same time. This recording technique is known as helical scan (see figure 7).

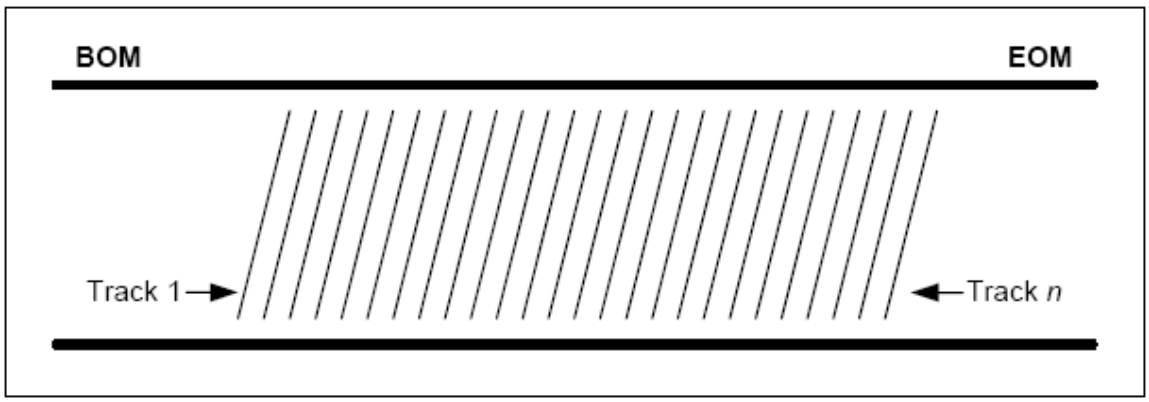

Figure 7 — Helical scan recording example

For most recording formats, a format identification in the form of a tone burst or some other recognizable pattern is recorded outside the user data area. The format identification is an attribute of a data volume used for interchange purposes and is defined in applicable standards.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

4.2.2.3 Cleaning medium model

As shown in figure Y, a portion of the physical length of a cleaning medium is not usable for abrading the read/write mechanism. For most cleaning volumes, a length of the cleaning medium is reserved between the take-up reel and the beginning-of-medium, and between the end-of-medium position and the supply reel. This is done to provide a sufficient tape wrap onto the reel hub and to ensure that interactions between the abrasive substrate and the read/write mechanism start in an undamaged section of the medium.

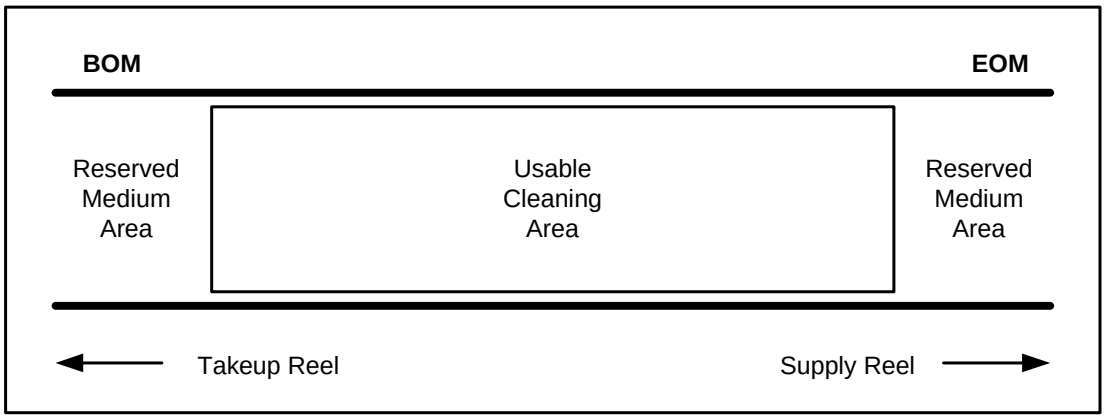

Figure Y — Typical cleaning medium layout

4.2.5 Partitions within a data volume

Partitions consist of one or more non-overlapped logical data volumes, each with its own beginning and ending points, contained within single physical data volume. Each partition (*x*) within a data volume has a defined beginning-of-partition (BOP *x*), an early-warning position (EW *x*), and an end-of-partition (EOP *x*).

All data volumes have a minimum of one partition called partition 0, the default data partition. For devices that support only one partition, the beginning-of-partition zero (BOP 0) may be equivalent to the beginning-of-medium and the end-of-partition zero (EOP 0) may be equivalent to the end-of-medium. For devices that support more than one partition, they shall be numbered sequentially starting with zero (i.e., beginning-of-partition 0).

When a data volume is mounted, it is logically positioned to the beginning of the default data partition (BOP 0). When a REWIND command is received in any partition (*x*), the device positions to the beginning-of-partition of the current partition (BOP *x*).

Partitions on a data volume may be recorded in any order and use any partition number unique to the physical data volume. It is sufficient for a device to be able to locate a partition, given its partition number, or to determine that it does or does not exist on the data volume. For interchange, information about which partitions are present on a data volume may be stored on the data volume in a format specified area, possibly unavailable to the application client, or the information may be an intrinsic attribute of the device implementation.

4.2.6.1 Logical objects within a partition

The basic unit of data transferred by an application client is called a logical block. Logical blocks are stored according to the specifications of the format for the data volume and may be recorded as portions of one or more physical blocks on the medium. The mapping between physical and logical blocks is the responsibility of the device server.

After writing data from BOP *x*, the medium is considered to be a contiguous grouping of logical objects. Depending on the format, blank medium may be treated as an end-of-data indication, an error recovery area, or an unrecoverable medium error causing an interchange error. Unrecorded data volumes, new or erased, may exhibit blank medium characteristics if an attempt is made to read or space the data volume before data has been written.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

4.2.10 Direction and position definitions

For sequential-access devices, positioning has the connotation of logically being in, at, before, or after some defined place within a data volume. Positioning requires that the position is capable of being repeated under the same circumstances. The orientation of usage for the four words (in, at, before, or after) is in one direction, from BOP *x* toward EOP *x*. All positioning defined below is worded from this perspective. Devices without object buffers have some physical position that relates to these logical positions. However, these definitions do not require the medium to have a physical position equivalent to the logical position unless explicitly stated.

The concept of being in some position means not being outside a defined region. The definition allows the position to be on the boundary of a defined region. When a data volume is first mounted, the logical position is always at the beginning of the default data partition (BOP 0). Whenever a data volume is mounted and the medium motion is stopped, the position is in some partition. While moving between partitions, there is no stable position.

4.2.12.1 Write protection introduction

Conditions such as positioning within unrecoverable data may result in a temporary write protection condition. To preserve future data integrity, the device server may reject any command that requires writing data to the medium when the recovery of the data is uncertain. A temporary write protection condition may be released by the device server at any time. Buffered logical objects may or may not be written to the media (e.g., the application client unloads the data volume before the temporary write protection condition is removed). The exact behavior of the device server during a temporary write protection condition is vendor specific.

Software write protection results when either the device server or medium is marked as write protected by a command from the application client. Four optional means of setting a software write protection state are available to an application client through the Device Configuration and Control mode pages:

a) software write protection for the device server across mounts;
b) associated write protection for the currently mounted data volume;
c) persistent write protection of a data volume across mounts; and
d) permanent write protection of a data volume across mounts.

4.2.12.4 Associated write protection

Associated write protection controls write protection for the currently mounted data volume as long as the current data volume is mounted. The associated write protection state is controlled by the ASOCWP bit in the Device Configuration mode page (see 8.3.3). Associated write protection exists if the ASOCWP bit is non-zero. Associated write protection may be altered by the application client (if the ASOCWP bit is changeable) if a data volume is mounted. If a data volume is de-mounted or after a logical unit reset occurs, associated write protection shall be removed.

4.2.12.5 Persistent write protection

Persistent write protection controls write protection for the currently mounted data volume. The persistent write protection state is controlled by the PERSWP bit in the Device Configuration mode page (see 8.3.3). If enabled, persistent write protection shall exist for the mounted data volume until disabled by the application client. The state of persistent write protection shall be recorded with the data volume and the persistent write protection shall only affect the application client accessible medium. The device server shall report the PERSWP bit as one when a mounted data volume is marked with persistent write protection. If a data volume is de-mounted or after a logical unit reset occurs, the device server shall report the PERSWP bit as zero prior to the mounting of a volume. The means for recording the state of persistent write protection for the data volume may be specified in the applicable recording format standard or be vendor specific.

4.2.12.6 Permanent write protection

Permanent write protection controls write protection for the currently mounted data volume. The permanent write protection state is controlled by the PRMWP bit in the Device Configuration mode page (see 8.3.3). If enabled, permanent write protection shall exist for the mounted data volume until disabled by a vendor-specific method. The state of permanent write protection shall be recorded with the data volume and the persistent write protection shall only affect the application client accessible medium. The device server shall report the PRMWP bit as one when a mounted data volume is marked with permanent write protection. If a data volume is de-mounted or after a logical unit reset occurs, the device server shall report the PRMWP bit as zero prior to the mounting of a volume. The means for recording the state of permanent write protection for the data volume may be specified in the applicable

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

recording format standard or be vendor specific. Permanent write protection shall not be removed by a MODE SELECT command using the PRMWP bit. Methods to remove this protection may or may not exist and are vendor specific.

4.2.16.1 TapeAlert introduction

Table 10 – TapeAlert flags

Flag	Name	Type	Severity	Deactivation condition
[Note: All rows not shown]				
0Bh	Cleaning ~~media~~ volume	O	I	Start of next medium load
14h	Cleaning operation required	O	C	After successful cleaning operation or cause resolved
15h	Cleaning operation requested	O	W	After successful cleaning operation
16h	Expired cleaning ~~media~~ volume	O	C	Start of next medium load
17h	Invalid cleaning ~~tape~~ volume	O	C	Start of next medium load

4.2.16.2.2 TapeAlert polling usage model

If using the TapeAlert polling usage model, the application client reads the TapeAlert log page or the TapeAlert Response log page without receiving notification from the device server that a TapeAlert flag has changed state. The application client may read the TapeAlert log page or the TapeAlert Response log page at any time (e.g. polled at a regular interval of 60 seconds). The application client should read either the TapeAlert log page or the TapeAlert Response log page:

a) ~~prior~~ prior to mounting a data volume and at the beginning of a data transfer sequence;
b) immediately after detecting an unrecoverable error during the data transfer sequence;
c) before de-mounting each data volume; and
d) at the end of a data transfer sequence.

4.2.20.2 Encrypting data on the medium

If data encryption is enabled for an I_T_L nexus and the mounted data volume supports the selected encryption algorithm at the current logical position, all logical blocks received by the device server from that I_T_L nexus as part of a WRITE(6) or WRITE(16) command shall be encrypted before being recorded on the medium. Filemarks are logical objects that shall not be encrypted.

If data encryption is enabled for an I_T_L nexus and the mounted data volume does not support the selected encryption algorithm at the current logical position, then the device server shall terminate a WRITE(6) or WRITE(16) command with CHECK CONDITION status, with the sense key set to DATA PROTECT, and the additional sense code set to ENCRYPTION PARAMETERS NOT USEABLE.

If data encryption is enabled for an I_T_L nexus and the mounted data volume does not support the selected encryption algorithm at the current logical position, then the device server may terminate a WRITE FILEMARKS(6) or WRITE FILEMARKS(16) command with CHECK CONDITION status, with the sense key set to DATA PROTECT, and the additional sense code set to ENCRYPTION PARAMETERS NOT USEABLE.

4.2.20.3 Reading encrypted blocks on the medium

A data volume may contain no encrypted blocks, all encrypted blocks, or a mixture of encrypted blocks and unencrypted blocks. The fact that logical blocks are encrypted shall not alter space or locate operations. The decryption mode shall be ignored when processing a filemark during a read or verify command.

4.2.20.4 Exhaustive-search attack prevention

If the device server has reached its limit on failed attempts to set the data encryption key or supplemental decryption keys and decrypt data, it shall disable decryption for all I_T nexuses. All subsequent SECURITY PROTOCOL OUT commands specifying the

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

Tape Data Encryption security protocol and with the SECURITY PROTOCOL SPECIFC field set to Set Data Encryption page with the DECRYPT field or ENCRYPT field set to any value other than DISABLE shall be terminated with CHECK CONDITION status, with the sense key set to DATA PROTECT, and the additional sense code set to DATA DECRYPTION KEY FAIL LIMIT REACHED. This condition shall persist until the data volume is de-mounted from the device or a hard reset condition occurs.

4.2.20.5 Keyless copy of encrypted data

In some scenarios it is desirable to copy data from one data volume to another without needing knowledge of the encryption parameters used to encrypt the data on the data volume.

A keyless copy logical unit (KCLU) controls configuration and data flows related to a data volume that is either a source or destination for encrypted data being transferred without requiring application client knowledge of an encryption key.

A keyless copy source logical unit (KCSLU) controls configuration and data flows related to the data volume from which the encrypted data is copied without requiring device server knowledge of an encryption key when the decryption mode is set to RAW.

A keyless copy destination logical unit (KCDLU) controls configuration and data flows related to the data volume to which the encrypted data is being copied without requiring device server knowledge of an encryption key when the encryption mode is set to EXTERNAL.

4.2.20.5 Managing keys within the device server

The device server shall release the resources used to save a set of data encryption parameters under the following conditions:

a) the CKOD bit is set to one in the saved data encryption parameters and the data volume is de-mounted;
b) the CKORL bit is set to one and the key scope is set to LOCAL in the saved data encryption parameters and the I_T nexus that established the set of data encryption parameters loses its reservation;
c) the CKORL bit is set to one and the key scope is set to ALL I_T NEXUS in the saved data encryption parameters and the device server experiences a reservation loss (see 3.1.55);
d) the CKORP bit is set to one in the saved data encryption parameters and the device server processes a PERSISTENT RESERVE OUT command with a service action of either PREEMPT or PREEMPT AND ABORT;
e) a microcode update is performed on the device; or
f) a power on condition occurs.

5.2 ERASE(16) command

Table 17 – method field values

Value	Description
00b	Vendor specific
01b	The device server shall erase or over-write the data volume with a format-specific pattern. Upon successful ~~procesiing~~ processing of the command, the data volume may contain fragments of data specified for erasure. The data specified for erasure shall not be recognizable as valid user data using normal data volume processing methods.
10b	The device server shall erase or over-write the data volume with a format-specific pattern(s). Upon successful processing of the command, the data volume shall not contain fragments of data specified for erasure.
11b	Reserved

If the Security Meta-Data (SMD) bit is set to one, the device server shall alter the security meta-data stored on the data volume with the method specified by the METHOD field. If the SMD bit is set to zero, the device server handling of the Security Meta-Data stored on the data volume is vendor specific.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

If the Vendor-specific Control Meta-data (VCM) bit is set to one, the device server shall alter the vendor-specific control meta-data stored on the data volume with the method specified by the METHOD field. If the VCM bit is set to zero, the device server handling of the vendor-specific control meta-data stored on the data volume is vendor specific.

7.1 FORMAT MEDIUM command

If the FORMAT field is 0h, the logical unit shall determine the format method to use. A valid FORMAT MEDIUM command with 0h in the FORMAT field shall cause all data on the entire physical data volume to be lost.

If the FORMAT field is 1h, the logical unit shall partition the medium using the current mode data from the Medium Partition mode page (see 8.3.4). If none of the mode bits SDP, FDP, or IDP are set to one, the device server shall return CHECK CONDITION. The sense key shall be set to ILLEGAL REQUEST with the addition sense code set to PARAMETER VALUE INVALID. If insufficient space exists on the medium for the requested partition sizes, the device server shall return CHECK CONDITION status. The sense key shall be set to MEDIUM ERROR and the additional sense code shall be set to VOLUME OVERFLOW. A valid FORMAT MEDIUM command with 1h in the FORMAT field may cause all data on the entire physical data volume to be lost.

If the FORMAT field is 2h, the logical unit shall perform the operations equivalent to a FORMAT field of 0h followed by a FORMAT field of 1h. A valid FORMAT MEDIUM command with 2h in the FORMAT field may cause all data on the entire physical data volume to be lost.

7.10 SET CAPACITY command

The SET CAPACITY command (see table 53) sets the available medium for the currently mounted data volume to a proportion of the total capacity of that volume. Any excess space shall be unavailable on the data volume after successful completion of this command until changed by a new SET CAPACITY command. This change shall persist through power cycles, logical unit resets, I_T nexus losses, and unloading or reloading of the data volume. Other vendor-specific actions such as physical erasure may change the total capacity of the data volume. The method for recording the available capacity and other marks needed to manage the resulting capacity for volume interchange may be specified in a recording format standard or may be vendor specific.

A valid SET CAPACITY command shall cause all data and partitioning information on the entire physical data volume to be lost. If the partitioning information changes, the device server shall generate a unit attention condition for all initiators with the additional sense code set to MODE PARAMETERS CHANGED.

The CAPACITY PROPORTION VALUE field specifies the portion of the total volume capacity to be made available for use. The CAPACITY PROPORTION VALUE field is the numerator to a fraction with a denominator of 65 535. The resulting available capacity on the data volume shall be equal to the total volume capacity multiplied by this fraction. The device server may round up the capacity to the next highest supported value. This rounding shall not be considered an error and shall not be reported.

> NOTE 47 Available and total volume capacities are approximate values that may be affected by defects that reduce the actual available capacity of the data volume. Other factors, such as partitioning, compression, and logical block packing may also affect available capacity.

8.2.2 Sequential-Access Device log page

The Sequential-Access Device log page defines data counters associated with data bytes transferred to and from the medium and to and from the application client, binary list parameters describing native capacities, and a binary list parameter related to cleaning operations.

A non-zero value of the cleaning requested parameter indicates that the device has requested a ~~head~~ cleaning operation and a subsequent cleaning ~~cycle~~ operation has not been completed. A zero value of the cleaning requested parameter indicates that the device has not requested a ~~head~~ cleaning operation. The cleaning requested parameter value shall persist across I_T nexus losses, logical unit resets, and power cycles.

8.2.5 Tape Diagnostic Data log page

The HOURS SINCE LAST CLEAN field contains the time in media motion (i.e., head) hours since the last successful cleaning operation at the time the command terminated with the CHECK CONDITION status. The HOURS SINCE LAST CLEAN field is equivalent to the value contained in the Device Statistics log page with a parameter code of 0008h at the time the command terminated with the CHECK CONDITION status.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

8.3.3 Device Configuration mode page

An associated write protection (ASOCWP) bit set to one specifies the logical unit shall inhibit all writing to the medium after performing a synchronize operation (see 4.2.12 and 4.2.12.4). When the ASOCWP bit is set to one, the currently mounted data volume is logically write protected until the data volume is de-mounted (see 4.2.12 and 4.2.12.4). When the ASOCWP bit is set to one, all commands requiring eventual writes to the medium shall return CHECK CONDITION status. The sense key shall be set to DATA PROTECT and the additional sense code should be set to ASSOCIATED WRITE PROTECT (see 4.2.12.2). An ASOCWP bit set to zero specifies the currently mounted volume is not write protected by the associated write protection. The ASOCWP bit shall be set to zero by the device server when the volume is de-mounted. This change of state shall not cause a unit attention condition. If the application client sets the ASOCWP bit to one while no volume is mounted, the device server shall terminate the MODE SELECT command with CHECK CONDITION status. The sense key shall be set to NOT READY and the additional sense code shall be set to MEDIUM NOT PRESENT. If the Device Configuration mode page is savable, the ASOCWP bit shall be saved as zero, regardless of the current setting.

A persistent write protection (PERSWP) bit set to one specifies the currently mounted data volume is logically write protected (see 4.2.12 and 4.2.12.5). When the PERSWP bit is set to one, all commands requiring eventual writes to the medium shall return CHECK CONDITION status. The sense key shall be set to DATA PROTECT and the additional sense code should be set to PERSISTENT WRITE PROTECT (see 4.2.12.2). A PERSWP bit set to zero specifies the currently mounted volume is not write protected by the persistent write protection. The PERSWP bit shall be set to zero by the device server when the volume is de-mounted or when a data volume is mounted with persistent write protection disabled. The PERSWP bit shall be set to one by the device server when a data volume is mounted with persistent write protection enabled. These changes of state shall not cause a unit attention condition. If the application client sets the PERSWP bit to one while no volume is mounted, the device server shall terminate the MODE SELECT command with CHECK CONDITION status. The sense key shall be set to NOT READY and the additional sense code shall be set to MEDIUM NOT PRESENT. If the application client sets the PERSWP bit to one when the logical position is not at BOP 0, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to POSITION PAST BEGINNING OF MEDIUM. If the Device Configuration mode page is savable, the PERSWP bit shall be saved as zero, regardless of the current setting.

A permanent write protection (PRMWP) bit set to one specifies the currently mounted data volume is logically write protected (see 4.2.12 and 4.2.12.6). When the PRMWP bit is set to one, all commands requiring eventual writes to the medium shall return CHECK CONDITION status. The sense key shall be set to DATA PROTECT and the additional sense code should be set to PERMANENT WRITE PROTECT (see 4.2.12.2). A PRMWP bit set to zero specifies the currently mounted volume is not write protected by the permanent write protection. The PRMWP bit shall be set to zero by the device server when the volume is de-mounted or when a data volume is mounted with permanent write protection disabled. The PRMWP bit shall be set to one by the device server when a data volume is mounted with permanent write protection enabled. These changes of state shall not cause a unit attention condition. If the application client sets the PRMWP bit to one while no volume is mounted, the device server shall terminate the MODE SELECT command with CHECK CONDITION status. The sense key shall be set to NOT READY and the additional sense code shall be set to MEDIUM NOT PRESENT. If the application client sets the PRMWP bit to one when the logical position is not at BOP 0, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to POSITION PAST BEGINNING OF MEDIUM. If the application client attempts to change the PRMWP bit from one to zero, the device server shall terminate the MODE SELECT command with CHECK CONDITION status. The sense key shall be set to DATA PROTECT and the additional sense code shall be set to PERMANENT WRITE PROTECT. If the Device Configuration mode page is savable, the PRMWP bit shall be saved as zero, regardless of the current setting.

8.3.4 Medium Partition mode page

The ADDITIONAL PARTITIONS DEFINED field specifies the number of additional partitions to be defined for a data volume when the SDP or IDP bit is set to one. The maximum value allowed is the value returned in the MAXIMUM ADDITIONAL PARTITIONS field. The ADDITIONAL PARTITIONS DEFINED value returned by the MODE SENSE command shall report one less than the number of partitions on the media when the logical unit is ready. If the unit is not ready, the ADDITIONAL PARTITIONS DEFINED field is undefined.

A logical unit is not required to retain the method used to partition the medium. The device server shall set only one of the IDP, FDP or SDP fields in the MODE SENSE data. If a data volume was previously partitioned through a MODE SELECT command with FDP or SDP set to one, a device server may set IDP to one in subsequent MODE SENSE data since the data volume has been initiator partitioned. However, in a MODE SELECT command, the application client cannot use IDP set to one in place of FDP or SDP set to one.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

The MEDIUM FORMAT RECOGNITION field specifies the logical unit's capability to automatically identify the medium format and partition information when reading a data volume. The value in this field may be different following a medium change. The MEDIUM FORMAT RECOGNITION field values are shown in table 82.

8.5.2.4 Data Encryption Capabilities page

The algorithm valid for mounted volume (AVFMV) bit shall be set to one if there is a data volume currently mounted in the device and the encryption algorithm being described is valid for that data volume. The AVFMV bit shall be set to zero if there is no volume mounted in the device or the algorithm is not valid for the currently mounted volume.

The distinguish encrypted data capable (DED_C) bit shall be set to one if the device server is capable of distinguishing encrypted data from unencrypted data when reading it from the medium. The DED_C bit shall be set to zero if the device server is not capable of distinguishing encrypted data from unencrypted data when reading it from the medium. If the ability to distinguish encrypted data from unencrypted data is format specific and a data volume is mounted, the DED_C bit shall be set based on the current format of the medium. If no volume is mounted, the DED_C bit shall be set to one if the device server is capable of distinguishing encrypted data from unencrypted data in any format that the device server supports.

The message authentication code capable (MAC_C) bit shall be set to one if the algorithm includes a message authentication code added to encrypted blocks. The MAC_C bit shall be set to zero if the algorithm does not include a message authentication code added to encrypted blocks. If the inclusion of a message authentication code is format specific and a data volume is mounted, the MAC_C bit shall be set based on the current format of the medium. If no volume is mounted, the MAC_C bit shall be set to one if the device server adds a message authentication code to data encrypted with this algorithm in any format that the device server supports.

The algorithm valid for current logical position (AVFCLP) field specifies if the encryption algorithm being specified is valid for writing to the mounted data volume at the current logical position. Table 102 specifies the values for the AVFCLP field.

Table 102 – AVFCLP field values

Code	Description
00b	Current logical position is not applicable to the encryption algorithm validity or no volume is loaded.
01b	The ~~ecryption~~ encryption algorithm being specified is not valid for writing to the mounted data volume at the current logical position.
10b	The ~~ecryption~~ encryption algorithm being specified is valid for writing to the mounted data volume at the current logical position.
11b	Reserved

8.5.3.2.1 Set Data Encryption page overview

If the clear key on de-mount (CKOD) bit is set to one the device server shall set the data encryption parameters to default values upon completion of a data volume de-mount. If the CKOD bit is set to zero, the de-mounting of a data volume shall not affect the data encryption parameters. If the CKOD bit is set to one and there is no volume mounted in the device, the device server shall terminate the command with CHECK CONDITION status and set the sense key to ILLEGAL REQUEST and the additional sense code to INVALID FIELD IN PARAMETER DATA.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304-1185
USA
www.hp.com

A.2 TapeAlert flag associated information

Table A.1 – TapeAlert log page parameter codes

Code	Flag	Recommended application client message	Probable cause
[Note: All rows not shown.]			
0Bh	Cleaning ~~media~~ volume	The tape in the drive is a cleaning cartridge.	Cleaning ~~tape~~ volume loaded into drive.
14h	Cleaning operation required	The tape drive needs cleaning: 1. If the operation has stopped, eject the tape and clean the drive. 2. If the operation has not stopped, wait for it to finish and then clean the drive. Check the tape drive user's manual for device specific cleaning instructions.	The drive thinks it has a head clog or needs a cleaning operation.
15h	Cleaning operation requested	The tape drive is due for routine cleaning: 1. Wait for the current operation to finish. 2. Then use a cleaning cartridge. Check the tape drive users manual for device specific cleaning instructions.	The drive is ready for a periodic cleaning operation.
16h	Expired cleaning ~~media~~ volume	The last cleaning cartridge used in the tape drive has worn out: 1. Discard the worn out cleaning cartridge. 2. Wait for the current operation to finish. 3. Then use a new cleaning cartridge.	The cleaning ~~tape~~ volume has expired.
17h	Invalid cleaning ~~tape~~ volume	The last cleaning cartridge used in the tape drive was an invalid type: 1. Do not use this cleaning cartridge in this drive. 2. Wait for the current operation to finish. 3. Then use a valid cleaning cartridge.	Invalid cleaning ~~tape~~ volume type used.