

**Working
Draft**

**T13
xxxxDT
Revision 0
June 14, 2004**

Information Technology - SMART Command Transport (SCT)

This is an internal technical report of T13, a Technical Committee of Accredited Standards Committee NCITS. As such, this is not a completed technical report and has not been approved. The contents may be modified by the T13 Technical Committee. The contents are actively being modified by T13. This technical report is made available for review and comment only.

Permission is granted to members of NCITS, its technical committees, and their associated task groups to reproduce this technical report for the purposes of NCITS standardization activities without further permission, provided this notice is included. All other rights are reserved. Any commercial or for-profit replication or republication is prohibited.

ASC T13 Technical Editor:

Curtis E. Stevens
Western Digital Technologies Inc.
20511 Lake Forest Dr.
Lake Forest, Ca. 92630

Tel: 949-672-7933
E-Mail: Curtis.Stevens@wdc.com

Reference number
ANSI TR INCITS.*** - 200x
Printed June, 17, 2004 9:14PM

T13/xxxxDT Revision 0

Other Points of Contact:

T13 Chair
Dan Colegrove
Hitachi Global Storage Technology Inc.
2505 Anthem Village Dr. Ste. E-221
Henderson, NV, 89052
Phone: 702-614-6119
Fax: 702-614-7955
E-Mail: Daniel.colegrove@hgst.com

T13 Vice-Chair
Jim Hatfield
Seagate Technology
389 Disc Drive
Longmont, CO 80503
Tel: 720-684-2120
Fax: 720-684-2711
E-mail: James.C.Hatfield@seagate.com

INCITS Secretariat
Administrator Standards Processing
1250 Eye Street, NW Suite 200
Washington, DC 20005
Tel: 202-737-8888
Fax: 202-638-4922
Email: INCITS@ITIC.ORG

T13 Reflector

See the T13 Web Site at <http://www.t13.org> for reflector information.

T13 Web Site

<http://www.t13.org>

T13 Anonymous FTP Site

<ftp.t13.org>

T13 mailings

Global Engineering
15 Inverness Way East
Englewood, CO 80112-5704
Tel: 303-792-2181 or 800-854-7179
Fax: 303-792-2192

Revision History		
Rev	Date	Description
0	June 14, 2004	Initial Release

American National Standard
for Information Systems —

SMART Command Transport (SCT)

Secretariat
Information Technology Industry Council

Approved mm dd yy

American National Standards Institute, Inc.

Abstract

This technical report describes a method for transporting commands to an ATA device using SMART log pages.

NCITS's Technical Report Series

This Technical Report is one of a series produced by the American National Standards Committee, NCITS, Information Technology. The secretariat for NCITS is held by Information Technology Industry Council (ITI), 1250 Eye Street, NW Suite 200, Washington DC 20005. As a by-product of the standards development process and the resources of knowledge devoted to it, NCITS from time to time produces Technical Reports. Such Technical Reports are not standards, nor are they intended to be used as such.

NCITS Technical Reports are produced in some cases to disseminate the technical and logical concepts reflected in standards already published or under development. In other cases, they derive from studies in areas where it is found premature to develop a standard due to still changing technology, or inappropriate to develop a rigorous standard due to the existence of a number of viable options, the choice of which depends on the users particular requirements. These Technical Reports, thus, provide guidelines, the use of which may result in greater consistency and coherence of information processing systems.

When the draft Technical Report is completed, the Technical Committee approval process is the same as for a draft standard. Processing by NCITS is also similar to that for a draft standard.

CAUTION: The developers of this Technical Report have requested that the holder's of patents that may be required for the implementation of this Technical Report, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this Technical Report.

As of the date of this Technical Report and the following calls for the identification of patents that may be required for the implementation of the Technical Report, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any Technical Report it Processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this Technical Report.

Published by
American National Standards Institute
11 West 42nd Street, New York, New York 10036

Copyright 200x by American National Standards Institute
All rights reserved.

Contents	Page
1 Foreword.....	iv
1 Scope.....	1
2 Definitions and abbreviations	2
2.1 Enhanced BIOS	2
2.2 Enhanced ATA Device.....	2
3 Command Transport.....	2
3.1 Capability Identification.....	2
3.2 Command Transport.....	3
3.3 Data Transfer3	
3.4 SCT Status Request.....	4
3.5 How to use SCT Pass-Through Commands	7
3.5.1 Issuing a SCT pass-thru command to the drive.....	7
3.6 Drive Access Methods	8
3.6.1 Logical Block Address.....	8
3.7 Command Decode Flow	9
3.8 SCT Command Nesting and intermingling with Standard commands	9
4 SCT Command Set.....	10
4.1 Long Sector Access command.....	11
4.2 LBA Segment Access command	12
4.3 Error Recovery Control command.....	12
4.4 Feature Control command.....	13
4.5 SCT Data Table command	15

Tables	Page
TABLE 1– IDENTIFY DEVICE WORD.....	2
TABLE 2– ISSUE SCT COMMAND	3
TABLE 2 - SUCCESSFUL SCT COMMAND RESPONSE	3
TABLE 3- SCT COMMAND ERROR RESPONSE	3
TABLE 4– READ\WRITE SCT DATA.....	4
TABLE 5– SCT STATUS REQUEST.....	4
TABLE 6 – FORMAT OF SCT STATUS RESPONSE	5
TABLE 7 – CONTROL COMMAND	7
TABLE 8– TRANSFER DATA TO HOST COMMAND	8
TABLE 9– TRANSFER DATA TO THE DRIVE COMMAND.....	8
TABLE 10 - KEY SECTOR FORMAT	10
TABLE 11 – SCT ACTION CODES.....	11
TABLE 12 – LONG SECTOR FORMAT	12
TABLE 13 – FEATURE CODE LIST	14
TABLE 14 – SCT DATA TABLES (BY TABLE IDENTIFIER).....	15
TABLE 15 – FAN CONTROL TEMPERATURE HISTORY TABLE FORMAT	16

1 Foreword

(This foreword is not part of American National Standard X3.***-199x.)

This technical report documents a method for passing commands that are not currently documented in the ATA/ATAPI series of standard to a device using log pages. These log pages can be accessed as a part of SMART, or they may be accessed as a part of the READ/WRITE LOG EXT commands. The commands that are executed by this mechanism include:

1. Long Sector Access (Read/Write Long)
2. LBA Segment Access (Fill sectors or copy sectors)
3. Error Recovery Control (Limits the time for error recovery in READ and WRITE commands)
4. Feature Control (overrides write cache enable/disable)
5. SCT Data Table command (returns HDA Temperature and Fan Ctrl Temperature history)

This technical report was developed by the ATA ad hoc working group of T13 during 2004. The approval process started in 2004.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the NCITS Secretariat, Information Technology Industry Council, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This technical report was processed and approved for submittal to ANSI by Accredited Standards Committee on Information Processing Systems, INCITS. Committee approval of the technical does not necessarily imply that all committee members voted for approval. At the time it approved this technical, the NCITS Committee had the following members:

[Editors Note: INCITS officers and members go here]

Technical committee T13 on ATA Storage Interfaces, which reviewed this standard, had the following members:

[Editors Note: T13 officers and members go here]

Other ad hoc participants were:

[Editors Note: Table of on-member participants goes here]

Introduction

This technical report encompasses the following:

Clause 1 describes the scope.

Clause 2 provides definitions, abbreviations, and conventions used within this technical report.

Clause 3 describes the transport mechanism.

Clause 4 describes the SCT command set

Clause 5 provides sample usages

AMERICAN NATIONAL STANDARD**INCITS.***-200x**

American National Standard
for Information Systems —

**Information Technology —
SMART Command Transport**

1 Scope

This technical report assumes that the reader is familiar with the ATA/ATAPI standard which include SMART, READ LOG EXT and WRITE LOG EXT.

ATA devices are increasing used in bridged environments where the Parallel ATA or Serial ATA interface is bridged to a different kind of bus. Sometimes this bridging requires that commands be processed or translated before passing them on to the ATA device. If the bridge is not designed to pass a specific command, then there is no way to issue that command to an ATA device. This technical report defines a way to pass commands to an ATA devices using SMART or extended logs. By passing the commands through a log write, the bridge sees a write log command. The ATA device can then parse the log entry and execute the intended command.

2 Definitions and abbreviations

2.1 Enhanced BIOS

All Enhanced BIOSs shall adhere to this technical report. Enhanced BIOSs should support at least one of the following:

- PIO Mode 3 or greater;
- DMA Mode 1 or greater.

2.2 Enhanced ATA Device

An Enhanced ATA Device is a hard disk or other device which interfaces to the system via Integrated Drive Electronics (IDE). These devices must conform to ATA-2(X3.279-1996)/ATAPI(SFF-8020 rev 2.6) or later and shall support:

- PIO Mode 3 or greater and;
- DMA Mode 1 or greater.

3 Command Transport

There are several phases involved in the issuing and executing of SCT commands. These phases are:

1. Capability Identification
2. Command Transport
3. Data Transfer
4. Status

Capability Identification is performed by issuing IDENTIFY DEVICE and checking to make sure the command is supported. Command Transport occurs when a 512 byte data packet is created and then written to SMART or extended log page E0h. The 512 byte data packet contains one or more commands as documented later in this report. If the command was a data transfer command then the data can be transferred by reading or writing log page E1h. If the data is larger than a single log page, the page is read or written multiple times. Finally, command status can be read at anytime by reading log page E0h. If the command involves data transfer, the host will need to check status before data is transferred to ensure that the drive is ready, as well as when the command is complete, to confirm that the data was transferred successfully. When the command is complete, the host may need to check status a third time to determine if the command succeeded, failed, or partially succeeded.

3.1 Capability Identification

The ATA IDENTIFY DEVICE command is used to determine if SCT is enabled and which SCT Action Codes are supported.

Table 1– IDENTIFY DEVICE Word

Word	O/M	F/V	Description
206	M		SCT Command set support.
		X	15 - 6 Reserved
		F	5 SCT Data Table (AC5) supported
		F	4 Features Control command (AC4) supported
		F	3 Error Recovery Control (AC3) supported
		F	2 LBA Segment Access (AC2) supported
		F	1 Long Sector Access (AC1) supported
		F	0 SCT command set supported (includes SCT status)

3.2 Command Transport

The command transport uses log sectors to pass-thru commands as well as inquire about status and control data flow. Two log pages have been allocated to support sending commands, receiving data, and checking status. Log page E0h is used for control and status. Log page E1h is used for reading or writing data. The command process is started by sending a “key” sector to log page E0h (the Control Port). The key sector informs the drive of the Action to be taken, any special mode of operation required for that action, and any other action-specific parameters.

Table 2– Issue SCT Command

Register	Value
Features	D6h (SMART Write Log subcommand code)
Sector Count	01h (Must be used for all pass-thru commands)
Sector Number	E0h (Must be used for all pass-thru commands)
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART op code)

Table 3 - Successful SCT Command Response

Register	Value
Status	50h
Error	00h
Sector Count	Depends on command (LSB)*
Sector Number	Depends on command (MSB)*
Cylinder	Number of sectors of data available

*For commands that return data in the task file (“Task File Output”), the data will be placed in the Sector Count and Sector Number registers

Table 4- SCT Command Error Response

Register	Value
Status	51h
Error	04h
Sector Count	Extended Error code, LSB*
Sector Number	Extended Error code, MSB*
Cylinder	Number of sectors of data available—up to and including the failed sector

*See Appendix A for list of Extended Error Codes

3.3 Data Transfer

Once an SCT command has been issued, status can be checked and data can be transferred. Data transfer uses log page E1h. Up to 255 sectors of data can be transferred at a time. If the SCT Command requires more than 255 sectors of data transfer, it can be written or read in 255 sector increments.

Table 5– Read\Write SCT Data

Register	Value
Features	D5h\D6h (Read data\Write data)
Sector Count	Cnt (Number of sectors to be transferred)
Sector Number	E1h
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART op code)

3.4 SCT Status Request

Once an SCT command has been issued, a status is reported in the ATA registers. This status indicates that the command was accepted or that an error occurred. This return does not indicate successful completion of the SCT actions. Some commands can take several minutes to execute. In this case, the host can determine execution progress by requesting SCT status. Some commands may need setup time before they are ready to receive data. SCT status is also used to determine when the device is read to receive data.

Log page E0 contains the status information. Reading log page E0h retrieves the status information. The SCT status may be acquired any time the task file is not busy. This command will not change the power state of the drive, nor terminate any background activity, including any SCT command in progress.

Table 6– SCT Status Request

Register	Value
Features	D5h (SMART Read Log subcommand code)
Sector Count	01h (Must for all pass-thru commands)
Sector Number	E0h (Must for all pass-thru commands)
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART op code)

Table 7 – Format of SCT Status Response

Offset (words)	Size (words)	Field Name	Description
0	1	Format Ver	Status Response format version number (=0x0001)
1	1	SCT Version	SCT command set version number MSB = major version (Function change) LSB = minor version (non-functional change/correction) Note, 0xYYZZ is displayed as yy.zz, where YY hex is in yy decimal 1.00 = Original release (Action code 1 to 3) 1.01 = AC4 – Feature Control, Disable/Enable Write Cache and “Zeroed Bit” support
2	1	SCT Spec.	Highest level of SCT specification supported
3	2	Status Flags	Bit 0 : Drive Zeroed. If this bit is set to 1, an LBA Segment Accesses command write to all LBA’s of the drive completed without error. This bit shall be cleared to 0 when any user LBA is written. It is also cleared if the capacity of the drive is change via HPA or DCO. This bit is persevered through a power cycle. Bits 1 to 31 : Reserved
5	1	Drive State	0 = Idle 1 = Stand-by 2 = Sleep 3 = DST executing in background 4 = SMART Off-line Data Collection executing in background 5 = SCT command executing in background
6	1	reserved	0x00
7	1	Error Code	Error of last SCT command. 0xFFFF if SCT command executing in background.
8	1	Action Code	Action code of last SCT command
9	1	Function Code	Function Code of last SCT command
10	10	reserved	Reserved for later use
20	4	LBA	Current LBA of SCT command executing in background
24	76	reserved	0x00
100	1	HDA Temp	Current drive HDA temperature in degrees Celsius
101	1	Max Temp	Maximum HDA temperature in degrees Celsius seen this power cycle
102	1	Life Max Temp	Maximum HDA temperature in degrees Celsius seen for the life of the drive

Offset (words)	Size (words)	Field Name	Description
103	1	Fan Ctrl Temp	<ul style="list-style-type: none"> • FFh (255) = Indicates invalid Fan Ctrl Temp value • 0F0h (240) = Over range (drive operating above allowable temperature limits) • 0D0h (208) = drive operating at max allowable temperature limits. (Typical max allowable value is 55C) <p>The following values are decimal:</p> <ul style="list-style-type: none"> • 5 = 5C under the drive operating allowable temperature • 10 = 10C under the drive operating allowable temperature • 15 = 15C under the drive operating allowable temperature • 20 = 20C under the drive operating allowable temperature • 25 = 25C under the drive operating allowable temperature • 30 = 30C under the drive operating allowable temperature • 35 = 35C under the drive operating allowable temperature • 40 = 40C under the drive operating allowable temperature • 45 = 45C under the drive operating allowable temperature • 50 = 50C under the drive operating allowable temperature • 55 = 55C under the drive operating allowable temperature • 60 = 60C under the drive operating allowable temperature • 65 = 65C under the drive operating allowable temperature <p>For example, if max allowable temperature is 55C and the current HDA Temp is 37C, Fan Ctrl Temp is computed by:</p> <p>55 – 37 = 18. Round up to next 5C increment.</p> <p>Yields: 20C or 14h</p> <p>The fan control ranges are product specific values selected to give optimal control of a fan that will provide airflow to the drive. The main purpose of having a value for every five degree Celsius delta is to allow the system to change the</p>

Offset (words)	Size (words)	Field Name	Description
104-495	-	reserved	
496-511	16	Vendor Specific	Vendor Specific

3.5 How to use SCT Pass-Through Commands

SCT commands are referred to as “pass-through” commands because they piggy-back on standard ATA commands. In this case, SMART Read Log and SMART Write Log commands or READ LOG EXT and WRITE LOG EXT. As viewed through the eyes of a standard ATA protocol analyzer, the SCT command will be seen as data being transferred by these commands; whereas from the perspective of a disk drive that implements this commands set, this “data” would be interpreted as either an SCT command request, an SCT command response, SCT status or SCT command data.

3.5.1 Issuing a SCT pass-thru command to the drive

Step 1 : Build the key sector in a host buffer for the appropriate action and parameters

Step 2 : Issue the control command (Table 8), and send the key sector to the drive.

Table 8 – Control Command

Register	Value
Features	D6h (SMART Write Log subcommand code)
Sector Count	01h (Must always be 01h)
Sector Number	E0h (Control and status port)
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART opcode)

Note, drive should respond with a 5000h status (see Table 3). If you get a command abort (5104h), then either the key sector format is invalid, the task file contains an invalid value or the command encountered an execution error. Check Sector Count, Sector Number registers for error code (see Table 4 and See **Error! Reference source not found.**). If the command was a “write” command, the command is terminated and will not accept any data; skip step 3. However, if the command was a “read” command, there maybe partial output available.(For example, on a sector read command, the data up to and including the sector in error, will be available.) In this case, the host may choose to proceed to step 3 to get the partial data. In certain cases, the error is not fatal and serves only as a warning.

If the status is 5000h, then check the Cylinder register. If Cylinder value 0, then command is complete and terminated without error. If Cylinder value > 0, then proceed to step 3.

Step 3 : Now use the Read Log Sector (Table 9) to transfer data from the drive to the host and use Write Log Sector (Table 10) command to transfer data from the host to the drive. The Sector Count register must be set to the number of sectors to be transferred. The transfer request must be in the range of 1 sector up to the total number of sectors not yet transferred. The number of sectors available will have been posted in the Cylinder register in the previous step. If the requested number of sectors is larger than the number of the sectors available, the drive will post an error. If the value is less than the number of sectors available, the Host may repeat Step 3 until all sectors have been transferred. The drive will auto advance the sector pointer by the number of sectors transferred, and report in the Cylinder register the number of sector remaining to be transferred (Cyl=0 means the command is done). In other words, the host has complete control over the number of sectors to transfer at a time. Note, if the number of sectors to be transferred is greater or equal to FFFFh, the drive will post FFFFh in the Cylinder Register and the value will remain at FFFFh until the number of sectors remaining drops below FFFFh. The exact number to be transferred is reported by the SCT Status command.

Step 4 : For key commands that require transfer of data to the drive (such as a write command), typically the command is not complete until the last block of data has been transferred to the drive. Upon receiving the last block of data, the drive will perform the specified operation (in this case, write the data to the disk). However, in the case of very large amounts of data, such as write all, some data may be processed (written to the disk) prior to receiving all of the data from the host.

Table 9– Transfer Data to Host Command

Register	Value
Features	D5h (SMART Read Log subcommand code)
Sector Count	Cnt (Number of sectors to transfer)
Sector Number	E1h (Data port)
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART opcode)

Table 10– Transfer Data to the Drive Command

Register	Value
Features	D6h (SMART Write Log subcommand code)
Sector Count	Cnt (Number of sectors to transfer)
Sector Number	E1h (Data port)
Cylinder	C24Fh (SMART enable code)
Command	B0h (SMART opcode)

3.6 Drive Access Methods

Standard ATA commands employ either LBA or Logical CHS access methods; SCT commands only support 48-bit LBA access.

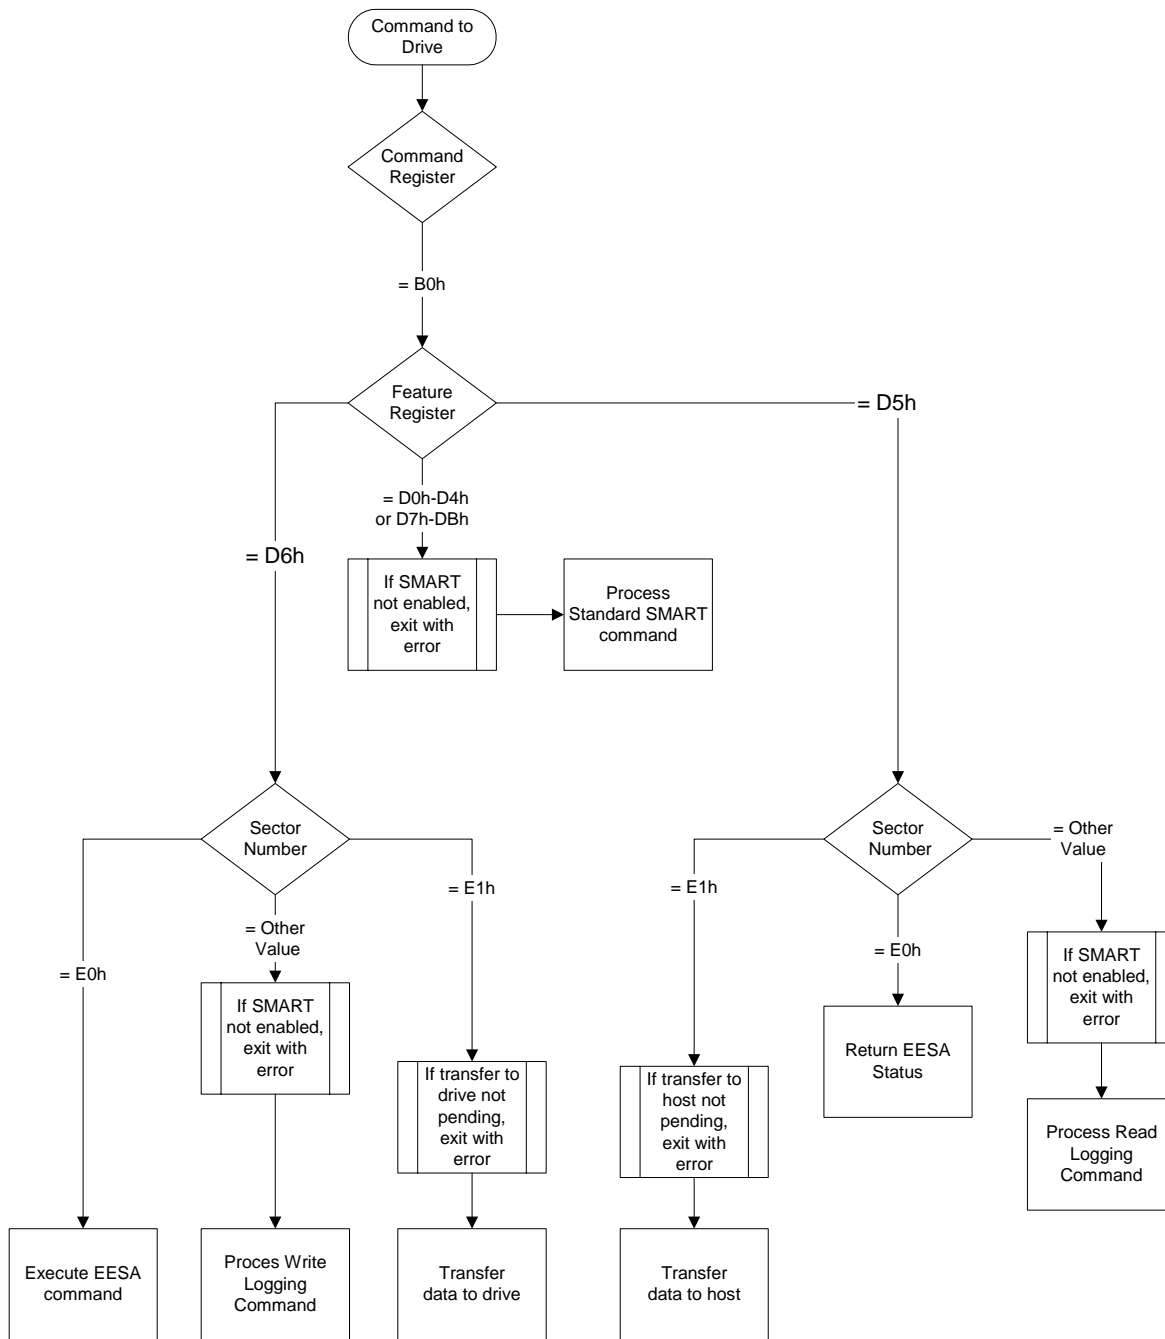
3.6.1 Logical Block Address

For LBA access all user sectors on the drive are numbered in a one-dimensional sequence from 0 to the maximum number of user sectors minus one. ATA Commands support 28-bits of LBA addressing and ATA Extended commands support 48-bits of LBA addressing. All SCT commands support 48-bits of LBA address. In

this method, all defective cylinders, heads and sectors are mapped out by defect management, rendering them inaccessible.

3.7 Command Decode Flow

The diagram below summarizes the command decode flow with respect to SCT operation.



Note, register values not defined in the flow will cause an error to be reported to the host.

3.8 SCT Command Nesting and intermingling with Standard commands

In general, standard ATA commands can be intermingled with SCT Commands but SCT commands can not be nested. SCT commands that do not require a follow-on data transfer operation never have an issue with being intermixed with any ATA commands or each other. SCT commands that do require a data transfer, on the other hand, may not be nested; that is, if a key command that requires a data transfer is issued all data transfer—to or from the host—must complete before another SCT command is issued. In most cases, however, ATA read/write commands may be inserted in between SCT data transfers, that is, between complete SMART Read Log/Write Log commands. Furthermore, any reset (POR, software or hardware) will cause the SCT command to be aborted.

4 SCT Command Set

An SCT command (a.k.a. as a Key Sector) is always 512 bytes long and takes the following format:

Table 11 - Key Sector format

Field	words	Description
Action Code	1	This is the fields that defines the command type and generally specifies the type of data being accessed, such as <i>sector</i> , <i>long sector</i> , etc. or physical action being performed, such as <i>seek</i> .
Function Code	1	This field specifies the type of access, such as <i>read</i> , <i>write</i> , <i>verify</i> , <i>set</i> , <i>clear</i> , etc
Parameter1	Depends on command	Depends on command
Parameter2	Depends on command	Depends on command
.	.	.
.	.	.
.	.	.
Total Words	256	

The supported action codes are defined below:

Table 12 – SCT Action Codes

Action Code	Block Data	TF Data	Description
0x0000	-	-	Reserved
0x0001	in\out	Y	Long Sector Access
0x0002	in	N	LBA Segment Access
0x0003	-	Y	Error Recovery Control
0x0004	-	Y	Features Control
0x0005	out	N	SCT Data Tables
0x0006 to 0xBFFF	-	-	Reserved
0xC000 to 0xFFFF	-	-	Vendor specific

4.1 Long Sector Access command

Action Code(1): **0x0001**

Function Code(1): 1 for read Long
 2 for write Long

LBA(4): Address of long sector (in LSW to MSW order)

Block Transfer: Output for read
 Input for write

Task File Output: Number of ECC/CRC bytes for read and write (reported in Sec Cnt and Sec Num registers)

Description: Read or write a sector with full CRC and ECC data

The Long Sector format for both read and write is two (512 byte) blocks long. The first block contains the user data; the second block contains the error correction and detections bytes; the remainder of the second block should be all zeros.

Table 13 – Long Sector Format

Field Name	Size (in Bytes)	Description
First Block		
User Data	512	User data as written from host
Second Block		
ECC/CRC Data	Depends on actual sector size	Error correction and detection bytes. The number of bytes is returned as TF Data on both read and write.
Reserved	Remainder of block	All zeros

4.2 LBA Segment Access command

<u>Action Code(1):</u>	0x0002
<u>Function Code(1):</u>	1 for repeat write pattern 2 for repeat write sector
<u>Start LBA(4):</u>	First LBA of segment (in LSW to MSW order)
<u>Sector Count(4):</u>	Number of sectors; if zero, do remainder of drive
<u>Pattern(2):</u>	32-bit pattern to be written for entire segment (FC 1 only)
<u>Block Transfer:</u>	One sector input for repeat write sector
<u>Task File Output:</u>	None
<u>Description:</u>	Write LBA segment

This command will begin writing sectors in LBA order until Sector Count sectors have been written. A Sector Count of zero means apply operation from Start LBA until the last user LBA on the drive. (Last LBA is determined by the HPA feature. This command will not write over a hidden partition is set by the Host Protected Area.)

Any command, other than SCT status, issue to the drive while this command is in progress will terminate the command.

Use the SCT Status command to determine status of the command, such as active, current LBA and any errors. While this command is in progress, the SCT status error code will be 0xFFFF, and set to 0x0000 if the command completes without error. It will be non-value (<0xFFFF) if the command terminated prematurely for any reason.

Drive Zeroed Flag. Issuing this command with a function code of 1 or 2 and a value of zero for Start LBA and Sector Count will cause all LBA's of the drive to be written with the specified pattern. If command completes without encountering an error or being aborted, then the "Drive Zeroed flag" (bit 0 of the Status Flags in the SCT status. See SCT status) will be set to 1. A sequential write to any LBA on the drive (except another complete write all), will cause the Drive Zeroed flag to be cleared.

4.3 Error Recovery Control command

<u>Action Code(1):</u>	0x0003
<u>Function Code(1):</u>	1 to set new value 2 to return current value

<u>Selection Code(1):</u>	1 for Read Command Timer (1 Count = 100ms) 2 for Write Command Timer (1 Count = 100ms)
<u>Set Value(1):</u>	New value to be set (Function Code 1 only)
<u>Block Transfer:</u>	None
<u>Task File Output:</u>	If Task File status is 5000h: Current value (Function Code 2 only)
<u>Description:</u>	Set error recovery time

This command can be used to set time limits for read and write error recovery. This may be desirable in a data redundant RAID environment where it is more desirable to have the drive report a data error rather than risk having it being kicked off of the RAID.

The Read Command Timer sets an upper limit to the amount of time the drive's disk task will be operating on a command. Typically amount of the time the drive will be operating on a read command in total but in some cases a read command will require more than one disk operation. Minimum value for this command is TBD. Setting this value to zero will disable Read Command time-out, allowing the drive to perform all available error recovery procedures without time limit.

If the Read Command Timer expires while the drive is performing error recovery, the drive will stop processing the command and report an un-correctable ECC error for the LBA that was causing error recovery to be invoked. Note that the LBA might actually be recoverable given more time for error recovery. At this point the host could reconstruct the data for the failing LBA from the other disk drives, and then issue a write command to the target LBA, allowing the drive to handle the suspect LBA as it best sees fit.

The Write Command Timer sets the upper limit for the amount of time the drive spends recovering from an error while performing a write. The minimum value for this command is TBD. Setting this value to zero will disable Write Command time-out, allowing the drive to perform all available error recovery procedures without a time limit.

The Write Command Timer has the effect of controlling how aggressively the drive will relocate write data when encountering servo defects. A long Write Command Timer allows the drive to use all of its available error recovery procedures for dealing with write errors. A short Write Command Timer will force the drive to relocate sectors that are having difficulty being written sooner rather than later. The relocating of the data needs to occur in order to prevent the write command from exceeding its allotted time.

The typical usage for an ATA or SATA drive is with Write Cache enabled. With write cache enabled, the drive cannot report an error on a write command. This is because the write command that is actually experiencing difficulty was probably finished from the host's perspective many commands ago. This leaves no recourse for the drive other than to relocate any sectors that are having difficulty being written.

Read and Write command timer values are set to default values at power-on but may be altered by a SCT command at any time. These settings are unaffected by software (soft) or hardware (pin 1) reset.

4.4 Feature Control command

<u>Action Code(1):</u>	0x0004
<u>Function Code(1):</u>	1 to set new state for selected feature 2 to return current state for selected feature 3 to return current Option Flags for selected feature

- Feature Code(1): Identifier of feature to be acted upon
- New State(1): Feature Code dependent value
- Option Flags(1): All options apply to all Features (except if state otherwise)
 Bit 0 – Set Permanent (SP). If this bit is set, the state change become the new default state and will be preserved through a power cycle. If cleared, the default set will be restore upon power cycle.
 Bit 1-15 – Reserved
- Block Transfer: None
- Task File Output: Current feature state (reported in Sec Cnt and Sec Num registers)
- Description: Set new feature state or return current state or Option Flags

Table 14 – Feature Code List

Feature Code	State Definition
0x0001	<p>0x0001 = Allow write cache operation to be determined by ATA Set Features command</p> <p>0x0002 = Force write cache enabled.</p> <p>0x0003 = Force write cache disabled</p> <p>If State 0x0001 is selected, the ATA Set Features command will determined the operation state of write cache per the ATA specification. If States 0x0002 or 0x0003 are selected, write cache will be forced into the corresponding operation state, regardless of the current ATA Set Features state. Furthermore, any future use of the ATA Set Features command shall be treated as a NOP, not effect the operation state of write cache and complete normally without reporting an error.</p> <p>In all cases, bit 5 of word 85 in the IDENTIFY DEVICE information will reflect the true operation state of write cache, one indicting enabled and zero indicating disabled.</p> <p>The default is states 0x01</p>
0x0002	<p>0x0001 = Enable Write Cache Reordering (use WRPO ordering for writes)</p> <p>0x0002 = Disable Write Cache Reordering (use FIFO ordering for writes)</p> <p>If State 0x0002 is selected, disk write scheduling is executed on a first-in-first-out (FIFO) basis, instead of Write Cache Rotational Position Optimization (WRPO) order. If State 0x0001 is selected, then RPO ordering is used for writes. If write cache is disable, the current Write Cache Reordering state is remembered but has no effect on non-cached writes, which are always written in the order received. The state of Write Cache Reordering has no effect on either NCQ or LCQ queued commands.</p> <p>The default is states 0x0001</p>

0x0003	<p>Set time Interval between entries for temperature Data Tables.</p> <p>0x0000 is invalid 0x0001 to 0xFFFF logging interval in minutes</p> <p>This value always applies equally to Fan Control Temperature History and HDA Temperature History queues. Issuing this command will cause the queues to be reset and any current entries to be lost: Queue Index shall be set to zero and the first queue location shall be set to current Fan Control / HDA Temperature and all remaining queue locations set to FFh</p> <p>Default value is 0x0001</p>
--------	--

4.5 SCT Data Table command

<u>Action Code(1):</u>	0x0005
<u>Function Code(1):</u>	1 to read data table
<u>Table Id(1):</u>	Identifier of Data Table to be acted upon
<u>Block Transfer:</u>	Number of sectors in table
<u>Task File Output:</u>	none
<u>Description:</u>	Access and control SCT Data Tables

Function Code 1 will return the specified data table..

Table 15 – SCT Data Tables (by Table Identifier)

Table Id	Description
0x0000	Invalid
0x0001	Fan Control Temperature History Table (in relative degrees C).
0x0002	HDA Temperature History Table (in absolute degrees C).
0x0003 to 0xCFFF	reserved
0xD000 to 0xFFFF	vendor specific

Table 16 – Fan Control Temperature History Table format

Offset (bytes)	Size (bytes)	Field Name	DESCRIPTION
0	2	Format Version	Data table format version (= 0x0001)
2	2	Sampling Period	The fan control sampling period in minutes. This is how often the drive actually samples its temperature sensor. 0x0000 is an invalid sampling period and indicates sampling is disabled.
4	2	Interval	This is the timer interval between entries in the history queue. By default, this value is 1 minute. This value shall be not be less than the Sampling Period.
6	1	Fan Max Op Limit	Fan control maximum operating temperature limit
7	1	Fan Over Limit	Fan Control over range temperature limit
8	22	reserved	Must be zero
30	2	Queue Size	Number of entry locations in history queue
32	2	Queue Index	Last updated entry in queue. Queue Index is zero-based, so Queue Index 0x0000 is the first location in the buffer (at offset 34) and Queue Index 0x007F is the last location in the buffer (at offset 161). See note 1.
34	128	Queue Buffer	History of the last 128 HDA Temperature values. This is a circular buffer so the newest value will overwrite the oldest value. A value of 0x00 indicates no value set. All samples are separated by Sampling Period in time. See note 1
162 to 255	-	reserved	Must be zero

Note 1 – Temperature History is preserved across power cycles, but with the requirement that when the drive powers up, the current entry in the queue will be set 0xFF, a invalid temp value. This way someone viewing the history could see that the discontinuity in temp result from the drive being turned off. If the drive does not sample HDA temperature during a certain power state—such as Sleep or Standby, then the same should be done when returning to a power state that does sample HDA temperature

Note 2 – When the Temperature history is cleared, for new drives and after changing Logging Interval, the Queue Index shall be set to zero and the first queue location shall be set to current Fan Control / HDA Temperature and all remaining queue locations set to FFh

Table 3 – HDA Temperature History Table format

Offset (bytes)	Size (bytes)	Field Name	Description
0	2	Format Version	Data table format version (= 0x0001)
2	2	Sampling Period	The fan control sampling period in minutes. This is how often the drive actually samples its temperature sensor. 0x0000 is an invalid sampling period and indicates sampling is disabled.
4	2	Interval	This is the timer interval between entries in the history queue. By default, this value is 1 minute. This value shall be not be less than the Sampling Period.
6	24	reserved	Must be zero
30	2	Queue Size	Number of entry locations in history queue
32	2	Queue Index	Last updated entry in queue. Queue Index is zero-based, so Queue Index 0x0000 is the first location in the buffer (at offset 34) and Queue Index 0x007F is the last location in the buffer (at offset 161). See note 1.
34	128	Queue Buffer	History of the last 128 HDA Temperature values This is a circular buffer so the newest value will overwrite the oldest value. A value of 0x00 indicates no value set. All samples are separated by Sampling Period in time. See note 1
162 to 255	-	reserved	Must be zero

Note 1 – Temperature History is preserved across power cycles, but with the requirement that when the drive powers up, the current entry in the queue will be set 0xFF, a invalid temp value. This way someone viewing the history could see that the discontinuity in temp result from the drive being turned off.

5 Usage Examples

[Editor's Note: Insert an example of each command]