

**Working
Draft**

**T13
1701DT-N**
Revision 3
October 29, 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

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@HitachiGST.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
1	August 16, 2004	<ol style="list-style-type: none"> 1. Changed from CHS registers to LBA Low/Mid/High 2. Updated formatting 3. Changed notation from 0x to xxxxh 4. Created an overview section 5. Added a section describing resets 6. Added a section describing functionality when various commands are supported/not supported vs enabled/disabled 7. Added General Purpose Logging command samples 8. Added Extended Error Codes 9. Updated Glossary to include TF DATA and Block Data 10. This document does not document the "Super Command" 11. Clarified the temperature sensor to be vendor specific
2	September 2, 2004	<ol style="list-style-type: none"> 1. Made changes requested during August 24-26 working group. 2. Added Security to flow chart in overview. 3. Changed flow chart to match normal flow chart conventions. 4. Removed examples in section 5.2 because the tables were duplicates 5. Changed command documentation to a tabular format in order to make things more clear. 6. Updated Absolute temperatures to be 2's complement. Also made the illegal value 80h 7. Updated Relative temperatures to be unsigned values. Also made the illegal value FFh 8. Clarified that all temperature values are 1 byte 9. Changed SCT Status (Table 11) to have only byte values for temperatures and limits 10. Added a figure to show the relationship between absolute and relative temperatures. 11. Changed step 3 to be more clear in talking about data transfer 12. Changed step for to be more clear in talking about command completion 13. Updated SCT command responses to show when a data transfer value is present. 14. Added back the "New State" field that got dropped from the Feature Control command. 15. Fixed the table 11 temperature formulas 16. Added a tolerance statement on the sampling period.
2a	October 1, 2004	<ol style="list-style-type: none"> 1. Updated figure 1 to be symmetrical and fix a few nits 2. Removed references to RPO in table 17 3. Changed numbers from decimal to hex in figure 2 4. Changes in/out references to read/write in table 15 5. Added verbiage in 5.1 to indicate that a read or write of log page E1 is expected after the key sector is accepted. 6. Reworded verbiage about the drive zeroed flag in 5.2 to make it clearer. Also added the same verbiage that was added for

		<p>5.1.</p> <ol style="list-style-type: none"> 7. Fixed wording issue in table 20 note 2. Cut and paste referenced relative temp instead of absolute temp. 8. Added invalid revision code, device is waiting for host to write data, and device is waiting for host to read data to table 11
<p>2b</p>	<p>October 5, 2004</p>	<ol style="list-style-type: none"> 1. Clarified Max Op Limit wording in table 19 to remove the ambiguity of the word between. 2. Modified table 19 Over Limit description to remove ambiguity from the word attained. 3. Upped the revision on Table 20 because the fill value was changed from FFh to 80h 4. After making the changes to table 19 and table 11 listed above to support relative temperature for fan control we decided to delete table 19 and change table 11 to not report fan control 5. Insert maximum and minimum operation values and limits in the absolute temperature table that is now table 19. 6. Changed the offsets in the inputs part of 5.2 to match the data types. 7. Changed the data type of the pattern to match the description. 8. In table 17 feature code 3, removed reference to relative temperature table. Also added full listing of upper and lower limits as items preserved. 9. Clarified wording in section 5.3 regarding the interaction of error timeouts with queued commands. 10. Fixed a cut and past error in note 1 of table 19.

ANSI®
INCITS.*-200x**

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
FOREWORD	III
INTRODUCTION	V
1 SCOPE	1
2 DEFINITIONS AND ABBREVIATIONS	1
2.1 BLOCK DATA	1
2.2 BYTE	1
2.3 DWORD	1
2.4 PASS-THROUGH COMMAND	1
2.5 QWORD	2
2.6 RESERVED	2
2.7 TF DATA	2
2.8 WORD	2
3 OVERVIEW	2
4 COMMAND TRANSPORT	4
4.1 CAPABILITY IDENTIFICATION	4
4.2 COMMAND TRANSPORT	4
4.3 DATA TRANSFER	7
4.4 SCT STATUS REQUEST	8
4.5 HOW TO USE SCT PASS-THROUGH COMMANDS	11
4.5.1 <i>Issuing an SCT pass-through command to the drive</i>	11
4.6 DRIVE ADDRESSING METHODS	12
4.6.1 <i>Logical Block Address</i>	12
4.7 SCT COMMAND NESTING AND INTERMINGLING WITH STANDARD COMMANDS	12
4.8 RESETS	13
4.9 SMART AND GENERAL PURPOSE LOGGING SUPPORT	13
5 SCT COMMAND SET	13
5.1 LONG SECTOR ACCESS COMMAND	14
5.2 LBA SEGMENT ACCESS COMMAND	15
5.3 ERROR RECOVERY CONTROL COMMAND	16
5.4 FEATURE CONTROL COMMAND	18
5.5 SCT DATA TABLE COMMAND	20

Tables	Page
Table 1 – IDENTIFY DEVICE Word	4
Table 2 – Issue SCT Command Using SMART	5
Table 3 – Issue SCT Command Using GPL.....	5
Table 4 – Successful SCT Command Response.....	5
Table 5 – SCT Command Error Response	6
Table 6 – Extended Status Codes.....	7
Table 7 – Read\Write SCT Data using SMART	8
Table 8 – Read\Write SCT Data using GPL.....	8
Table 9 – SCT Status Request Using SMART.....	9
Table 10 – SCT Status Request Using GPL	9
Table 11 – Format of SCT Status Response	10
Table 12 – SCT Command Using SMART.....	11
Table 13 – SCT Command Using GPL	11
Table 14 – Key Sector Format.....	13
Table 15 – SCT Action Codes	14
Table 16 – Long Sector Format.....	15
Table 17 – Feature Code List.....	19
Table 18 – SCT Data Tables (by Table Identifier).....	20
Table 19 – Absolute HDA Temperature	21
Figures	Page
Figure 1– Command Decode	3

Foreword

(This foreword is not part of American National Standard INCITS.***-200x.)

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 Relative and Absolute HDA Temperature and 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]

T13/1701DT Revision 3

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 provides an overview

Clause 4 describes the transport mechanism.

Clause 5 describes the SCT command set

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 Block Data

Block Data is the data transferred to or from the device using SCT read/write log capabilities.

2.2 Byte


A byte is a unit of data that consists of eight bits as described below:



2.3 DWord

A DWord (Double Word) is a unit of data that consists of four bytes. This data may be represented on paper as a series of bits numbered from 31 to 0. In memory byte 0 of a DWord is stored in the lowest byte address and byte 3 is stored in the highest byte address.

On Paper:



In Memory:




2.4 Pass-Through Command

SCT commands are referred to as “pass-through” commands because they piggy-back on standard ATA command.


2.5 QWord

A QWord (Quad Word) is a unit of data that consists of eight bytes. This data may be represented on paper as a series of bits numbered from 63 to 0. In memory byte 0 of a QWord is stored in the lowest byte address and byte 7 is stored in the highest byte address.

On Paper:



In Memory:



2.6 Reserved

Reserved is a keyword indicating reserved bits, bytes, words, fields, and code values that are set aside for future standardization. Their use and interpretation may be specified by future extensions to this or other standards. A reserved bit, byte, word, or field shall be set to zero. The recipient should not check reserved bits, bytes, words, or fields. Receipt of reserved code values in defined fields should be treated as an error.


2.7 TF Data

TF Data is data returned in the ATA register block (TaskFile Data) in response to an SCT command.


2.8 Word

A Word is a unit of data that consists of two bytes. This data may be represented on paper as a series of bits numbered from 15 to 0. In memory byte 0 of a Word is stored in the lower byte address and byte 1 is stored in the higher byte address.

On Paper:



In Memory:



3 Overview

The basis of this technical report for issuing commands to a drive is log pages E0h and E1h. These log pages are used as follows:

	Log page E0	Log page E1
Write log page	Issue Command	Send data
Read log page	Return Status	Receive data

There are two ways to access the log pages: using SMART READ/WRITE LOG and READ/WRITE LOG EXT. Both sets of commands access the same log pages and provide the same capabilities. They are also used in the same way: A command is issued, data is transferred (if necessary), and status is retrieved.

The commands that are defined in SCT are subject to requirements documented in ATA/ATAPI-7. This means that if security is enabled and a password has not been issued to unlock the device, that media access commands will fail.

Figure 1 is a sample flowchart that shows how to process SCT log page requests using SMART:

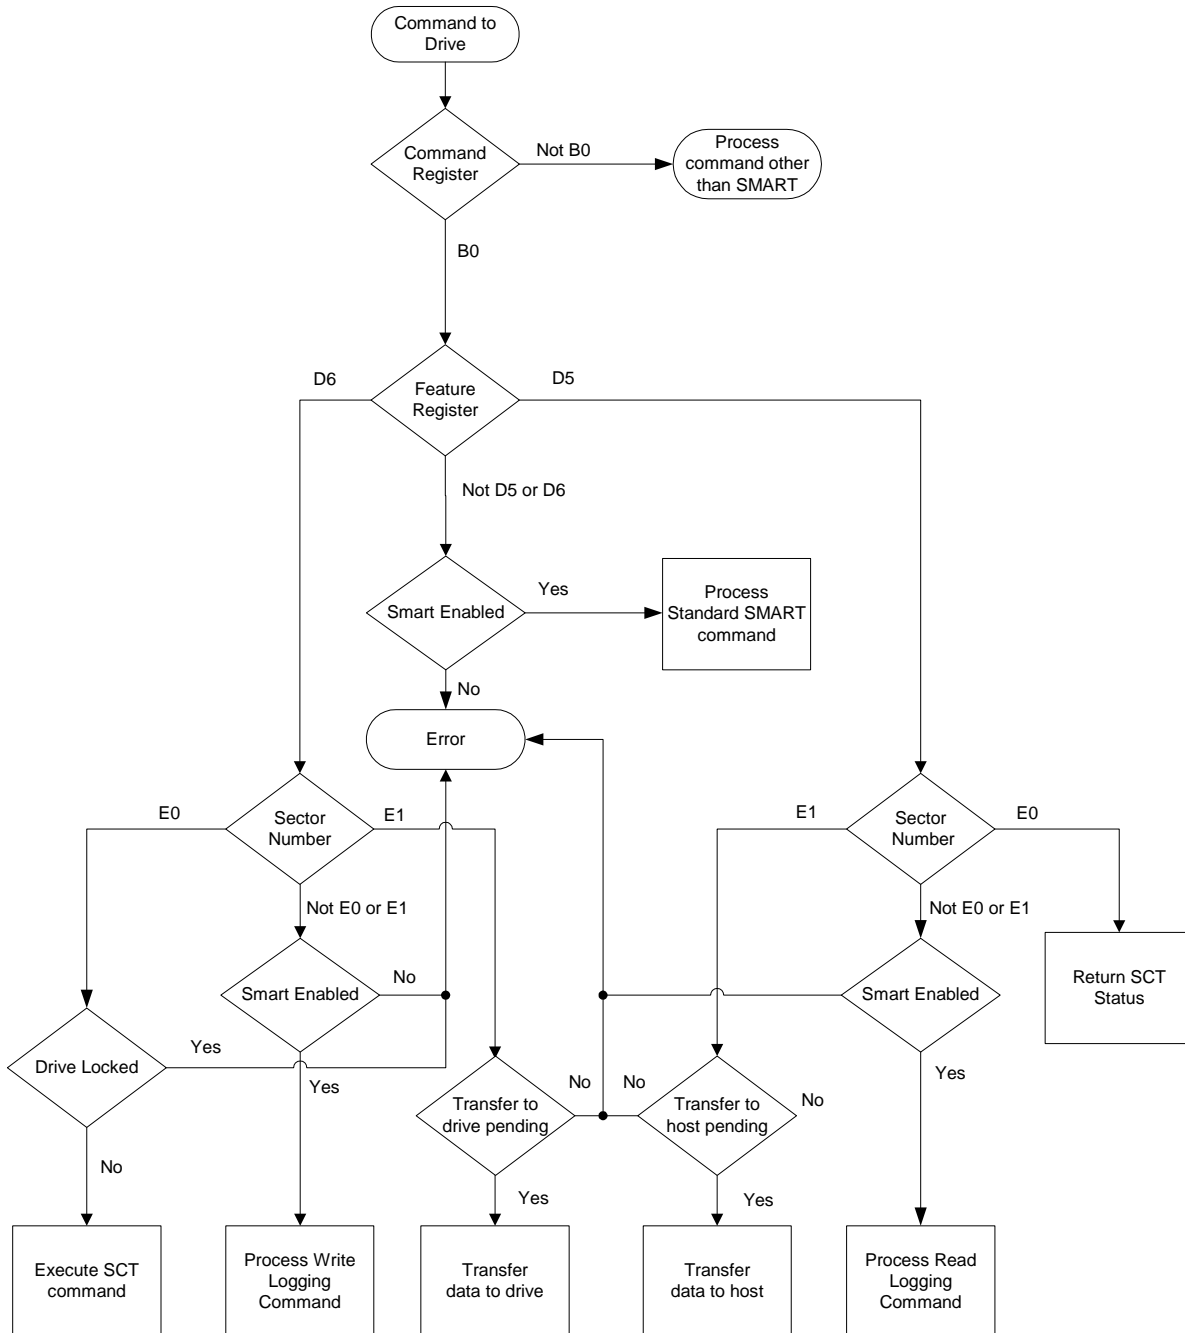


Figure 1– Command Decode

4 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 reading or writing log page E1h can transfer the data. 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.

4.1 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)

4.2 Command Transport

The command transport uses log sectors to pass-through commands as well as to 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 and writing data. Sending a "key" sector to log page E0h starts the command process. 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. The log pages can be accessed by SMART READ/WRITE LOG (See Table 2), or by the General Purpose Logging commands (GPL) READ LOG EXT and WRITE LOG EXT (See Table 3). The command successful response (See Table 4) and the error response (See Table 5) are the same for both the SMART and the GPL methods of issuing command.

Table 2 – Issue SCT Command Using SMART

Register	Value
Features	D6h (SMART Write Log subcommand code)
Sector Count	01h (Must be used for all pass-through commands)
LBA Low	E0h (Must be used for all pass-through commands)
LBA Mid	4Fh (SMART Enable Code)
LBA High	C2h (SMART Enable Code)
Command	B0h (SMART)

Table 3 – Issue SCT Command Using GPL

Register	Value
Features	Current Previous Reserved Reserved
Sector Count	Current Previous 01h 1 sector for pass-through commands 00h
LBA Low	Current Previous E0h Reserved
LBA Mid	Current Previous 00h There is no offset when commands are issued 00h
LBA High	Current Previous Reserved Reserved
Command	3Fh (WRITE LOG EXT)

Table 4 – Successful SCT Command Response

Register	Value
Error	00h
Sector Count	Depends on command (LSB)
LBA Low	Depends on command (MSB)
LBA Mid	Low 8 bits of number of sectors remaining to transfer for both read and write operations.
LBA High	High 8 bits of number of sectors remaining to transfer for both read and write operations.
Status	50h

Table 5 – SCT Command Error Response

Register	Value
Error	04h
Sector Count	Extended Status code, LSB (See Table 6)
LBA Low	Extended Status code, MSB (See Table 6)
LBA Mid	Low 8 bits of number of sectors of data available - up to and including the failed sector
LBA High	High 8 bits of number of sectors of data available - up to and including the failed sector
Status	51h

Table 6 – Extended Status Codes

Error Code	Definition
0000h	Command complete without error
0001h	Invalid Function code in LBA Segment Access command
0002h	Input LBA out of range
0003h	Request sector count overflow. The number of sectors requested to transfer (Sector Count register) in the read or write log command is larger than required by SCT command
0004h	Invalid Function code in Error Recovery command
0005h	Invalid Selection code in Error Recovery command
0006h	Host read command timer is less than minimum value
0007h	Host write command timer is less than minimum value
0008h	Background SCT command was aborted interrupting host command
0009h	Background SCT command was terminated because of unrecoverable servo error
000Ah	Invalid Function code in Long Sector Access command
000Bh	SCT data transfer command (SMART Read Log or SMART Write Log) was issued without first issuing SCT control command
000Ch	Invalid Function code in Feature Control command
000Dh	Invalid Feature code in Feature Control command
000Eh	Invalid New State value in Feature Control command
000Fh	Invalid Option Flags value in Feature Control command
0010h	Invalid SCT Action code
0011h	Invalid Table ID (table not supported)
0012h	Command was aborted due to drive security being enabled
0013h	Invalid revision code
0014h	Device is waiting for host to write data
0015h	Device is waiting for host to read data
0013h-BFFFh	Reserved
C000h-FFEFh	Vendor specific
FFF0h-FFFEh	Reserved
FFFFh	SCT command executing in background

4.3 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 and SMART logging is used to transfer the data, the data can be written or read in 255 sector increments. If GPL is used to transfer data, up to 65,535 sectors (approximately 33MB) can be transferred by a single command. If more than 65,535 sectors are required, then multiple GPL read or write logs

can be issued. Table 7 shows how to initialize the registers for data transfer using SMART. Table 8 shows how to initialize the registers for data transfer using GPL.

Table 7 – Read\Write SCT Data using SMART

Register	Value
Features	D5h\D6h (Read data\Write data)
Sector Count	Count (Number of sectors to be transferred)
LBA Low	E1h
LBA Mid	4Fh (SMART enable code)
LBA High	C2h (SMART enable code)
Command	B0h (SMART)

Table 8 – Read\Write SCT Data using GPL

Register	Value
Features	Current Previous Reserved Reserved
Sector Count	Current Previous Count (7:0) Count (15:8)
LBA Low	Current Previous E1h Reserved
LBA Mid	Current Previous 00h There is no offset when data is retrieved 00h
LBA High	Current Previous Reserved Reserved
Command	2Fh (READ LOG EXT)/3Fh (WRITE LOG EXT)

4.4 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 or even hours 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 BSY=0 and DRDY=1 and DRQ=0. This command will not change the power state of the drive, nor terminate any background activity, including any SCT command in progress. This means if the drive is in standby or idle the log request will still succeed. Table 9 shows how to initialize the registers for retrieving status using SMART. Table 10 shows how to initialize the registers for retrieving status using GPL. Table 11 shows the format of the status return information.

Table 9 – SCT Status Request Using SMART

Register	Value
Features	D5h (SMART Read Log subcommand code)
Sector Count	01h (Must be used for all pass-through commands)
LBA Low	E0h (Must be used for all pass-through commands)
LBA Mid	4Fh (SMART enable code)
LBA High	C2h (SMART enable code)
Command	B0h (SMART op code)

Table 10 – SCT Status Request Using GPL

Register		Value
Features	Current	Reserved
	Previous	Reserved
Sector Count	Current	01h 1 sector for pass-through status
	Previous	00h
LBA Low	Current	E1h
	Previous	Reserved
LBA Mid	Current	00h There is no offset when data is retrieved
	Previous	00h
LBA High	Current	Reserved
	Previous	Reserved
Command		2Fh (READ LOG EXT)

Table 11 – Format of SCT Status Response

Byte	Type	Field Name	Description
0	Word	Format Version	0002h - Status Response format version number.
2	Word	SCT Version	Manufacturer's vendor specific implementation version number
4	Word	SCT Spec.	0001h - Highest level of SCT Technical Report supported
6	DWord	Status Flags	Bit 0: Drive Zeroed. If this bit is set to 1, an LBA Segment Access command write to all LBAs of the drive has completed without error. This bit shall be cleared to 0 when any user LBA is written. This bit is also cleared if the capacity of the drive is change via SETMAX, SETMAX EXT or DCO. This bit is preserved through a power cycle. Bits 1-31: Reserved
10	Byte	Drive State	0 = Active waiting for a command 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
11	Byte	reserved	
12	Word	reserved	
14	Word	Extended Status Code	Status of last SCT command issued. FFFFh if SCT command executing in background (See Table 6).
16	Word	Action Code	Action code of last SCT command issued. If the Extended Status Code is FFFFh this is the Action Code of the command that is currently executing
18	Word	Function Code	Function code of last SCT command issued. If the Extended Status Code is FFFFh this is the Function Code of the command that is currently executing
20	Byte [20]	reserved	Reserved for later use
40	QWord	LBA	Current LBA of SCT command executing in background. If there is no command currently executing in the background, this field is undefined.
48	Byte [152]	reserved	00h
200	Byte	HDA Temp	Current drive HDA temperature in degrees Celsius. This is a 2's complement number. 80h indicates that this value is invalid.
201	Byte	Reserved	
202	Byte	Max Temp	Maximum HDA temperature in degrees Celsius seen this power cycle. This is a 2's complement number. 80h indicates that this value is invalid.
203	Byte	Reserved	
204	Byte	Life Max Temp	Maximum HDA temperature in degrees Celsius seen for the life of the drive. This is a 2's complement number. 80h indicates that this value is invalid.

Byte	Type	Field Name	Description
205	Byte	Reserved	
206	Byte [274]	reserved	Shall be cleared to zero
480	Byte [32]	Vendor Specific	

4.5 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 command set, this “data” would be interpreted as either an SCT command request, an SCT command response, SCT status or SCT command data.

4.5.1 Issuing an SCT pass-through command to the drive

4.5.1.1 Step 1 - Build a Key Sector

Build the key sector in a host buffer for the appropriate action and parameters.

4.5.1.2 Step 2 - Issue the SCT command

Issue the SCT command (Table 12 or Table 13), and send the key sector to the drive.

Table 12 – SCT Command Using SMART

Register	Value
Features	D6h (SMART Write Log subcommand code)
Sector Count	01h (Must always be 01h)
LBA Low	E0h (Command port)
LBA Mid	4Fh (SMART enable code)
LBA High	C2h (SMART enable code)
Command	B0h (SMART)

Table 13 – SCT Command Using GPL

Register	Value
Features	Current Previous Reserved Reserved
Sector Count	Current Previous 01h 1 sector for pass-through commands 00h
LBA Low	Current Previous E0h Reserved
LBA Mid	Current Previous 00h There is no offset when commands are issued 00h
LBA High	Current Previous Reserved Reserved
Command	3Fh (WRITE LOG EXT)

The device should respond with a 50h status (see Table 4). If there is a command abort (Status = 51h and Error = 04h), then either the key sector format is invalid, the task file contains an invalid value or the command

encountered an execution error. Check the Sector Count and LBA Low registers for the error code (see Table 5 and Table 6). 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 50h, then check the LBA Mid and LBA High registers. If the values are 0, then the command is complete and terminated without error; proceed to Step 4. If the values are greater than 0, then proceed to Step 3.

4.5.1.3 Step 3 - Transfer Data if Required

Read the SCT status response (See Table 9, Table 10, and Table 11) to determine that the device is ready for data transfer. To transfer data from the drive to the host issue a read log to page E1h (See Table 7 and Table 8) command. To transfer data from the host to the drive issue a write log to page E1h (See Table 7 and Table 8). 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 remaining will have been posted in the LBA Mid and LBA High registers in the previous step. If the requested number of sectors is larger than the number of the sectors remaining, the drive will post an error. If the value is less then the number of sectors remaining, the Host may repeat Step 3 until all sectors have been transferred. For commands that access the media, the drive will auto advance the sector pointer by the number of sectors transferred, and report in the LBA Mid and LBA High registers the number of sectors remaining to be transferred (if both registers are zero then the command is done, proceed to Step 4). 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 LBA Mid and High registers 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. Upon receiving the last block of data, the drive will perform the specified operation. In the case of very large amounts of data, such as LBA Segment Access, some data may be processed (written to the disk) prior to receiving all of the data from the host.

4.5.1.4 Step 4 - Final Status/SCT Command Completion

Read the SCT status response (See Table 9, Table 10, and Table 11) to determine how the command completed. If the command has not completed (by reporting FFFFh in Table 11 byte 14) then wait for some period of time and repeat Step 4 until the command is complete.

For SCT 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.

4.6 Drive Addressing Methods

Standard ATA commands employ either LBA or Logical CHS addressing using both 28-bit and 48-bit capability. SCT commands only support 48-bit addressing.

4.6.1 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.

4.7 SCT Command Nesting and intermingling with Standard commands

In general, standard ATA commands can be intermingled with SCT Commands but SCT commands cannot 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 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.

4.8 Resets

If an SCT command is executing, any reset including Soft Reset (SRST), Hard Reset, and Power-On Reset (POR) all cause the command to be terminated. This could result in partial command execution or data loss. There is no indication once the drive becomes ready that the previous command was terminated.

4.9 SMART and General Purpose Logging Support

If READ LOG EXT or WRITE LOG EXT is not supported, then the commands described in this technical report cannot be issued using the general purpose logging commands.

If SMART is not supported, then the commands described in this technical report cannot be issued using the SMART READ/WRITE LOG commands.

If SMART is supported, but not enabled, devices that implement this technical report still support SMART READ/WRITE LOG. This deviates from ATA/ATAPI-7 that states an error is returned for SMART READ/WRITE LOG when SMART is disabled.

5 SCT Command Set

An SCT command (Key Sector) is always 512 bytes long. Table 14 shows the generic format of an SCT command.

Table 14 – Key Sector Format

Field	Words	Description
Action Code	1	This field 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 and varies by command. For example, this can specify read, write, verify, etc.
Parameter1	Depends on command	Depends on command
Parameter2	Depends on command	Depends on command
.	.	.
.	.	.
.	.	.
Total Words	256	

The action codes are defined in Table 15.

Table 15 – SCT Action Codes

Action Code	Block Data	TF Data	Description
0000h	-	-	Reserved
0001h	Read/Write	Y	Long Sector Access
0002h	Write	N	LBA Segment Access
0003h	-	Y	Error Recovery Control
0004h	-	Y	Features Control
0005h	Read	N	SCT Data Tables
0006h - BFFFh	-	-	Reserved
C000h - FFFFh	-	-	Vendor specific

5.1 Long Sector Access command

Inputs:

Word	Name	Value	Description
0	Action Code	0001h	Read or Write a sector with full ECC or CRC data. The function of this action is similar to the ATA READ/WRITE LONG capability.
1	Function Code	0001h	Read Long
		0002h	Write Long
2	LBA	QWord	Sector to be read or written

Outputs:

Field Name	Value
Error	See ATA/ATAPI-7
Sector Count	Low Order number of ECC/CRC bytes
LBA Low	High Order number of ECC/CRC bytes
LBA Mid	Low 8 bits of number of sectors requested (Normally this will be two)
LBA High	High 8 bits of number of sectors requested (Normally this will be zero)
Device	Undefined
Status	See ATA/ATAPI-7

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 detection bytes; the remainder of the second block should be all zeros. Once the key sector has been issued and the TF Data indicates that the drive is ready to transfer data, log page E1h should be read or written to transfer the data.

Table 16 – Long Sector Format

Field	Size	Description
First Block		
User Data	512	This is the data normally sent or returned by a read or write command. This data may be encoded.
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

5.2 LBA Segment Access command

Inputs:

Word	Name	Value	Description
0	Action Code	0002h	This action writes a pattern or sector of data repeatedly to the media. This capability could also be referred to as "Write All" or "Write Same."
1	Function Code	0001h	Repeat Write Pattern
		0002h	Repeat Write Sector
2	Start	QWord	First LBA
6	Count	QWord	Number of sectors to fill
10	Pattern	DWord	If the Function Code is 0001h, this field contains a 32-bit pattern that is written on the media starting at the location specified in words two through five.

Outputs:

Field Name	Value
Error	See ATA/ATAPI-7
Sector Count	Undefined
LBA Low	Undefined
LBA Mid	Low 8 bits of number of sectors requested (Normally this will be one)
LBA High	High 8 bits of number of sectors requested (Normally this will be zero)
Device	Undefined
Status	See ATA/ATAPI-7

This command will begin writing sectors in incrementing order until Count sectors have been written. A Count of zero means apply operation from Start until the last user LBA on the drive is reached. The HPA feature determines the last user LBA. This command will not write over a hidden partition when hidden partitions are enabled using the Host Protected Area drive capabilities. Automatic sector reassignment is permitted during the operation of this function.

If Start or Start+Count go beyond the last user LBA then an error is reported and the SCT command is not executed. Issuing this command with a value of zero for Start LBA and Sector Count will cause all LBAs of the drive to be written with the specified pattern.

Any command, other than SCT status, issued to the drive while this command is in progress will terminate the command, including IDENTIFY DEVICE.

Use the SCT Status command to retrieve status information about the current SCT command. Example status information includes: command active or complete, current LBA, and errors. While this command is in progress, the SCT status error code will be FFFFh, and set to 0000h if the command completes without error. It will be less than FFFFh and greater than 0000h if the command terminated prematurely for any reason.

This command can change the Drive Zeroed Flag. If the command writes all the user addressable sectors and completes without encountering an error or being aborted, then the "Drive Zeroed flag" (bit 0 of the Status Flags in the SCT status. See Table 11) will be set to 1. Although this is called the Drive Zeroed Flag, it is also set when a non-zero pattern is written to all the user addressable LBAs. A write to any sector on the drive (except another complete write all), will cause the Drive Zeroed flag to be cleared.

Once the key sector has been issued, if the Function Code was 0002h and the TF Data indicates that the drive is ready to receive data, log page E1h should be written to transfer the data.

5.3 Error Recovery Control command

Inputs:

Word	Name	Value	Description
0	Action Code	0003h	Set the read and write error recovery time
1	Function Code	0001h	Set New Value
		0002h	Return Current Value
2	Selection Code	0001h	Read Timer
		0002h	Write Timer
3	Value	Word	If the function code is 0001h then this field contains the recovery time limit in 100ms units. A value of 1=100ms, 2=200ms, etc.

Outputs:

Field Name	Value
Error	See ATA/ATAPI-7
Sector Count	If Function Code was 0002h, then this is the LSB of the requested recovery limit. Otherwise, this field is undefined.
LBA Low	If Function Code was 0002h, then this is the MSB of the requested recovery limit. Otherwise, this field is undefined.
LBA Mid	Undefined
LBA High	Undefined
Device	Undefined
Status	See ATA/ATAPI-7

This command can be used to set time limits for read and write error recovery. For non-queued commands, these timers apply to command completion at the host interface. For queued commands where in order data delivery is enabled, these timers begin counting when the device begins to execute the command, not when the command is sent to the device. These timers do not apply to queued commands when out-of-order data delivery is enabled.

Time limits for error recovery 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. This is typically the amount of 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 one. 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 one. 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 reallocate 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 reallocating of the data needs to occur in order to prevent the write command from exceeding its allotted time. If the timer expires during a retry, the relocation operation is completed. If the timer is about to expire, it is the responsibility of the drive to relocate the data before the timer expires.

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 or COMRESET) reset.

5.4 Feature Control command

Inputs:

Word	Name	Value	Description						
0	Action Code	0004h	Set or return the state of drive features described in Table 17						
1	Function Code	0001h	Set new state for a feature						
		0002h	Return the current state of a feature						
		0003h	Return feature option flags						
2	Feature Code	Word	See Table 17 for a list of the feature codes						
3	New State	Word	Feature Code dependent value						
4	Option Flags	Word	<table border="0"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td>If the function code is 0001h, setting bit 0 to one causes the requested feature state change to be permanent across power cycles. If the function code is 0001h, setting bit 0 to zero causes the requested feature state change to be volatile. A hard reset causes the drive to revert to default, or last non-volatile setting.</td> </tr> </tbody> </table>	Bit	Description	15:1	Reserved	0	If the function code is 0001h, setting bit 0 to one causes the requested feature state change to be permanent across power cycles. If the function code is 0001h, setting bit 0 to zero causes the requested feature state change to be volatile. A hard reset causes the drive to revert to default, or last non-volatile setting.
Bit	Description								
15:1	Reserved								
0	If the function code is 0001h, setting bit 0 to one causes the requested feature state change to be permanent across power cycles. If the function code is 0001h, setting bit 0 to zero causes the requested feature state change to be volatile. A hard reset causes the drive to revert to default, or last non-volatile setting.								

Outputs:

Field Name	Value
Error	See ATA/ATAPI-7
Sector Count	If the Function Code was 0002h this is the Feature State LSB If the Function Code was 0003h this is the Option Flags LSB Otherwise this field is undefined
LBA Low	If the Function Code was 0002h this is the Feature State MSB If the Function Code was 0003h this is the Option Flags MSB Otherwise this field is undefined
LBA Mid	Undefined
LBA High	Undefined
Device	Undefined
Status	See ATA/ATAPI-7

Table 17 – Feature Code List

Feature Code	State Definition
0001h	<p>0001h = Allow write cache operation to be determined by ATA Set Features command</p> <p>0002h = Force write cache enabled.</p> <p>0003h = Force write cache disabled</p> <p>If State 0001h is selected, the ATA Set Features command will determined the operation state of write cache per the ATA specification. If State 0002h or 0003h is selected, write cache will be forced into the corresponding operation state, regardless of the current ATA Set Features state. Any attempt to change the write cache settings through SET FEATURES shall be treated as a NOP, and not affect 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 indicating enabled and zero indicating disabled.</p> <p>The default is states 0001h</p>
0002h	<p>0001h = Enable Write Cache Reordering</p> <p>0002h = Disable Write Cache Reordering</p> <p>If State 0002h is selected, disk write scheduling is executed on a first-in-first-out (FIFO) basis. If State 0001h is selected, then disk write scheduling may be reordered by the drive. If write cache is disabled, 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 state 0001h</p>
0003h	<p>Set time Interval between entries for temperature Data Tables.</p> <p>0000h is invalid</p> <p>0001h to FFFFh logging interval in minutes</p> <p>This value applies to the Absolute HDA Temperature History queue. Issuing this command will cause the queue to be reset and any current entries to be lost. Queue Index shall be set to zero and the first queue location for will be set to the proper current value. All remaining queue locations are set to 80h. The Sample Period, Max Op Limit, Over Limit, Min Op Limit and Under Limit values are preserved.</p> <p>Default value is 0001h</p>

5.5 SCT Data Table command

Inputs:

Word	Name	Value	Description
0	Action Code	0005h	Read a data table
1	Function Code	0001h	Read Table
2	Table ID	Word	See Table 18 for a list of data tables

Outputs:

Field Name	Value
Error	See ATA/ATAPI-7
Sector Count	Undefined
LBA Low	Undefined
LBA Mid	Low 8 bits of number of sectors remaining (Normally this will be one)
LBA High	Low 8 bits of number of sectors remaining (Normally this will be zero)
Device	Undefined
Status	See ATA/ATAPI-7

The Placement, accuracy, and granularity of temperature sensors to support Table 19 are vendor specific.

Table 18 – SCT Data Tables (by Table Identifier)

Table Id	Description
0000h	Invalid
0001h	Reserved
0002h	HDA Temperature History Table (in absolute degrees C). See Table 19
0003h to CFFFh	Reserved
D000h to FFFFh	Vendor Specific

Table 19 – Absolute HDA Temperature

Byte	Size	Field Name	Description
0	Word	Format Version	0002h - Data table format version
2	Word	Sampling Period	Absolute HDA Temperature sampling period in minutes. This is how often the drive actually samples its temperature sensor. 0000h is an invalid sampling period and indicates sampling is disabled.
4	Word	Interval	This is the timer interval between entries in the history queue. By default, this value is 1 minute. This value should not be less than the Sampling Period.
6	Byte	Max Op Limit	Maximum recommended continuous operating temperature. See Note 3. This is a one byte 2's complement number, that allows a range from -127°C to +127°C to be specified. 80h is an invalid value. This is a fixed value.
7	Byte	Over Limit	Maximum temperature limit. Operating the device above this temperature may cause physical damage to the device. See Note 3. This is a one byte 2's complement number, that allows a range from -127°C to +127°C to be specified. 80h is an invalid value. This is a fixed value.
8	Byte	Min Op Limit	Minimum recommended continuous operating limit. See Note 3. This is a one byte 2's complement number, that allows a range from -127°C to +127°C to be specified. 80h is an invalid value. This is a fixed value.
9	Byte	Under Limit	Minimum temperature limit. Operating the device below this temperature may cause physical damage to the device. See Note 3. This is a one byte 2's complement number, that allows a range from -127°C to +127°C to be specified. 80h is an invalid value. This is a fixed value.
10	Byte [20]	reserved	Shall be zero
30	Word	Queue Size	Number of entry locations in history queue. This number shall be in the range of 128-478.
32	Word	Queue Index	Last updated entry in queue. Queue Index is zero-based, so Queue Index 0000h is the first location in the buffer (at offset 34). The last entry in the buffer is at Queue Index + Queue Size (at Byte 34 + Queue Index + Queue Size). See Notes 1 and 2.

Byte	Size	Field Name	Description
34	Byte [Queue Size]	Queue Buffer	History of the last Queue Size Absolute HDA Temperature values. Other drive activities, such as data transfer, take priority over making this queue non-volatile. These are one byte 2's complement numbers, which allow a range from -127°C to +127°C to be specified. The Queue Buffer is a circular buffer of bytes. The newest value will overwrite the oldest value. A value of 80h indicates no value set. The Sampling Period separates all samples in time. See Note 1.
Queue Size + 34	Byte [512-Queue Size-34]	reserved	Shall be zero.
<p>Note 1 - The Absolute HDA Temperature History is preserved across power cycles with the requirement that when the drive powers up, the current entry in the queue shall be set to 80h, an invalid absolute temperature value. This way an application viewing the history can see the discontinuity in temperature result from the drive being turned off. If the drive does not sample temperatures 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.</p> <p>Note 2 - When the Absolute HDA Temperature history is cleared, for new drives and after changing the Logging Interval, the Queue Index shall be set to zero and the first queue location shall be set to the current Absolute HDA Temperature value. All remaining queue locations are set to 80h.</p> <p>Note 3 - These values should take into account the accuracy of the temperate sensor.</p>			