

Small Form Factor Committee
Specification Draft for
S.M.A.R.T. Applications Guide for the ATA and
SCSI Interfaces
SFF-8055i

June 5, 1996
Revision 1.4

This document was written by Fujitsu Ltd., Hewlett-Packard Company, IBM Corporation, Maxtor Corporation, Quantum Corporation, Seagate Technology, and Western Digital Corporation. It is endorsed by the S.M.A.R.T. Working Group. The S.M.A.R.T. Working Group consists of Compaq Computer Corporation, Dell Computer Corporation, Fujitsu Ltd., Hewlett-Packard Company, Hitachi Ltd., IBM Corporation, Maxtor Corporation, Micro House International, Micropolis (USA) Incorporated, NexStor, Phoenix Technologies Ltd., Promise Technology, Quantum Corporation, Seagate Technology, Symantec Corporation and Toshiba Corporation, Western Digital Corporation.

This document is provided by the S.M.A.R.T. Working Group "AS IS" and no maintenance or support services are provided by Fujitsu Ltd., Hewlett-Packard Company, IBM Corporation, Maxtor Corporation, Quantum Corporation, Seagate Technology, or Western Digital Corporation. Fujitsu, Hewlett-Packard, IBM, Maxtor, Quantum, Seagate, and Western Digital expressly disclaim all warranties, either express or implied, including but not limited to any implied warranties of merchantability, fitness for a particular purpose or against infringement of any copyright, patent, trade secret or other intellectual property right of any third party. You agree that Fujitsu, Hewlett-Packard, IBM, Maxtor, Quantum, Seagate, and Western Digital are not liable for any damages whatsoever, whether direct, indirect, incidental, special or consequential (including but not limited to loss of profits, revenue, data or use) incurred by you or any third party in connection with or arising out of the use of this document, even if Fujitsu, Hewlett-Packard, IBM, Maxtor, Quantum, Seagate, and Western Digital have been advised of the possibility of such damages.

"Self-Monitoring, Analysis and Reporting Technology", "S.M.A.R.T." and "The S.M.A.R.T. System" are not trademarks nor are they intended to be trademarks but rather are industry terms of art that describe technology used to monitor and predict device performance and/or reliability.

Document Status

Revision 1.0	-	Initial released document.
Revision 1.1	-	Draft update for inclusion of EXECUTE OFF-LINE IMMEDIATE
Revision 1.2	-	Clarification of EXECUTE OFF-LINE IMMEDIATE. Changed
Revision 1.3	-	Draft update for inclusion of SCS EXECUTE OFF-LINE IMMEDIATE
Revision 1.4	-	Draft update for minor changes requested during document review.

Table of Contents

1. INTRODUCTION	5
2. REFERENCES	5
3. SUPPORTED S.M.A.R.T. FUNCTIONS.....	6
3.1. SUPPORTED ATA FUNCTIONS.....	6
3.1.1. Enable/Disable Attribute Autosave	6
3.1.2. Execute Off-Line Immediate (Optional).....	6
3.1.3. Enable S.M.A.R.T. Operations.....	6
3.1.4. Disable S.M.A.R.T. Operations	7
3.1.5. Return S.M.A.R.T. Status	7
3.2. SUPPORTED SCSI FUNCTIONS.....	7
4. IMPLEMENTATION GUIDELINES	7
4.1. SYSTEM SOFTWARE.....	7
4.2. OPERATING SYSTEM CONSIDERATIONS.....	9
4.2.1. Windows Operating System Considerations for ATA Drives.....	9
4.2.1.1. Windows 3.1	9
4.2.2.1. Windows 95 and Windows NT.....	10
5. HARD DISK DRIVE IMPLEMENTATION GUIDELINES.....	10
5.1. ATA HARD DISK DRIVE GUIDELINES	10
5.2. SCSI HARD DISK DRIVE GUIDELINES.....	10
6. SOFTWARE APPLICATIONS GUIDE COMPLIANCE.....	10
<u>APPENDIX</u>	12

1. INTRODUCTION

The S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) working group was formed in order to guide the industry-wide implementation of S.M.A.R.T. with the intention of ensuring compatibility and customer satisfaction. This document is intended to be a guide for software applications, BIOS and driver development. S.M.A.R.T. is a technology developed to manage the reliability of data storage devices. S.M.A.R.T.-capable PC systems have the goal of enhancing system reliability by warning users of some pending device failures. With sufficient warning, users may have the opportunity to back up vital data and replace suspect devices prior to data loss or unscheduled down time. S.M.A.R.T. capability is a key new element in the PC architecture that will one day provide new levels of data integrity and data availability.

Peripheral data storage devices (such as hard disk drives) are complex electro-mechanical devices and, as such, can suffer performance degradation or failure due to a single event or a combination of events. Some events are immediate and catastrophic while others cause a gradual degradation of the device's ability to perform. It is possible to predict a portion of the failures, but S.M.A.R.T. cannot and will not predict all future device failures. S.M.A.R.T. should be treated as a feature to assist the computer user in preventing some but not all system down time due to device failure.

S.M.A.R.T. capable devices monitor a wealth of information internal to the device to assess reliability and predict an impending device failure. This information is, in some cases, available through the interface and can be presented to end-users via drivers and supporting applications. This data should not be presented to or interpreted by system users or managers to predict the integrity or reliability of a S.M.A.R.T. device. The predictive algorithms in a S.M.A.R.T. device are designed to interpret internal conditions in order to detect impending failures and thus users or system managers should not attempt to predict impending device failure from this internal data. S.M.A.R.T. data are not linear predictors of the degrading reliability of a S.M.A.R.T. capable device. It is the responsibility of a S.M.A.R.T. device to predict an impending failure and report that failure via the ATA interface RETURN S.M.A.R.T. STATUS command or the SCSI interface Informational Exception Condition.

2. REFERENCES

- Working Draft of the proposed American National Standard X3T10/2008D Revision 6 or higher, Information Technology AT Attachment-3 Interface (ATA-3).
- Working Draft of the proposed American National Standard X3T10/995D Revision 9 or higher, SCSI-3 Primary Commands.
- Compaq/Microsoft "Windows95 and Windows NT S.M.A.R.T. IOCTL API Specification", Revision 01.10 or higher.

3. SUPPORTED S.M.A.R.T. FUNCTIONS

3.1. SUPPORTED ATA FUNCTIONS

Table 1 lists those S.M.A.R.T. sub-commands and their respective sub-command codes recommended by the S.M.A.R.T. Working Group.

Table 1-- S.M.A.R.T. Sub-commands and Codes

Code	Sub-command
D2h	ENABLE/DISABLE ATTRIBUTE AUTOSAVE
D4h	EXECUTE OFF-LINE IMMEDIATE (optional)
D8h	ENABLE S.M.A.R.T. OPERATIONS
D9h	DISABLE S.M.A.R.T. OPERATIONS
DAh	RETURN S.M.A.R.T. STATUS

Future implementations of the S.M.A.R.T. feature may not implement all sub-commands listed in ATA-3.

3.1.1. Enable/Disable Attribute Autosave

Sub-command ENABLE/DISABLE ATTRIBUTE AUTOSAVE is intended to be used only to Enable the Attribute Autosave function during the application initialization. See section 4.1.

3.1.2. Execute Off-Line Immediate (Optional)

Sub-command EXECUTE OFF-LINE IMMEDIATE is optional and may be used with devices that support this vendor specific function. Collection of Attribute data in an “off-line” mode may have varying impact on device performance if the device is required to respond to commands from the host while performing its off-line data collection routine.

Note that for this command to be effectively used, status and information bytes need to be obtained from the data buffer returned on the S.M.A.R.T. Read Attribute Values command. See Appendix for details of this sub-command implementation.

3.1.3. Enable S.M.A.R.T. Operations

Sub-command ENABLE S.M.A.R.T. OPERATIONS instructs the device to start monitoring and predicting its health status.

3.1.4. Disable S.M.A.R.T. Operations

Sub-command DISABLE S.M.A.R.T. OPERATIONS instructs the device to cease all monitoring and predicting of its health status.

3.1.5. Return S.M.A.R.T. Status

Sub-command RETURN S.M.A.R.T. STATUS is used to determine if the device has a positive or negative health status.

3.2. SUPPORTED SCSI FUNCTIONS

Table 2 lists the S.M.A.R.T. reporting method and respective Method of Reporting Informational Exception (MIRE) code required by the S.M.A.R.T. Working Group.

Table 2-- S.M.A.R.T. Method of Reporting Informational Exception

Code	Method
04h	Unconditionally Generate Recovered Error

When negative S.M.A.R.T. health status is detected, the Unconditionally Generate Recovered Error reporting method requires the drive to present a check condition . The sense key for the check condition is Recovered Error and the additional sense code indicates an informational exception condition (see Section 2, SCSI-3 Primary Commands reference).

4. IMPLEMENTATION GUIDELINES

4.1. SYSTEM SOFTWARE

Software application, BIOS and direct port driver developers should adhere to the guidelines that follow.

BIOS and driver developers should ensure that S.M.A.R.T. and attribute autosaving are enabled immediately after device initialization. For ATA drives this requires an ENABLE S.M.A.R.T. command followed by an ENABLE ATTRIBUTE AUTOSAVE command. It is recommended that S.M.A.R.T. applications SHOULD NOT provide a means to explicitly enable or disable the Autosave feature. SCSI devices are enabled with a MODE SELECT command to page 1Ch, setting the Dexcept and Perf bits to 0. It is recommended that the Mode Page 1Ch Interval Timer field be set to 10 minutes to prevent performance degradation due to frequent negative health status notification.

Software applications should be tolerant of devices that do not support the S.M.A.R.T. feature. ATA devices that do not support the feature shall abort any S.M.A.R.T. commands received. The response of legacy ATA devices to S.M.A.R.T. commands may

be unpredictable. SCSI devices that do not support the S.M.A.R.T. feature do not have Mode Page 1Ch or do not support the Dexcept bit as a changeable bit. The application should be written to tolerate devices that both do and do not support the S.M.A.R.T. feature in the same system.

Due to vendor specific definitions, no conclusions can be drawn from reading attribute or threshold values. APPLICATIONS SHOULD NOT INTERPRET, OR PRESENT TO USERS, THE ATTRIBUTES OR THRESHOLD DATA VALUES.

S.M.A.R.T. applications shall prominently display a message that S.M.A.R.T. is only an advisory service and not an accurate predictor of device reliability.

Software applications shall be designed with the ability to indicate when a negative health status is received as a result of an ATA RETURN S.M.A.R.T. STATUS command or a SCSI Informational Exception Condition. The warning message should be displayed until it is physically acknowledged by the end user or reported as appropriate in more complex systems. After the event is acknowledged the application may optionally display an icon on the screen reminding the user or system manager that a failure prediction event has occurred.

When a negative status is received from the device, applications developers should use the following message that the S.M.A.R.T. Working Group endorses:

WARNING: Immediately back-up your data and replace your hard disk drive. A failure may be imminent.

The application shall indicate which physical device caused the S.M.A.R.T. warning message to be displayed. Additionally, the application should offer help that would include a list of the logical drives that correspond to the physical device returning the negative health status along with other useful information such as the type, manufacturer and serial number of the device.

In addition to informing the user of a predicted failure S.M.A.R.T. applications should:

1. Identify which devices on the system are S.M.A.R.T. capable and which are not.
2. Allow the user to individually enable or disable S.M.A.R.T. monitoring on each device (the default state for the application is enabled).
3. Identify which devices are S.M.A.R.T. enabled and which are disabled.
4. Poll ATA devices using the RETURN S.M.A.R.T. STATUS command to detect negative health status. Check SCSI devices for negative health status after each command that generates a check condition. Polling ATA devices less than once in 4 hours or issuing commands to SCSI devices less than once in 4 hours may significantly reduce the effectiveness of S.M.A.R.T.
5. Software Applications may disable SCSI S.M.A.R.T. after a negative health status has been detected to maximize performance during backup operations.

The following S.M.A.R.T. supplemental comments may be included in a Help or About window:

- The degrading condition may be caused by external conditions such as excessive heat or humidity. No matter what the cause, however, you should perform a system back up of your data NOW.
- S.M.A.R.T. stands for Self-Monitoring, Analysis and Reporting Technology. It monitors and reports operating degradation or fault conditions in the peripheral storage device that could lead to data loss with continued operation.
- S.M.A.R.T. capable devices monitor a wealth of information to assess reliability and predict an impending device failure. However, S.M.A.R.T. will not and cannot detect all impending failures. S.M.A.R.T. should be treated as a feature to assist the computer user in preventing some but not all system down time due to peripheral storage device failure.
- Data integrity can only be ensured by employing timely back-up procedures. Users should not delay back-up or reliability management procedures because of the presence of S.M.A.R.T. on their system.

The application developer may want to include the ability to display vendor specific messages. The actual messages may be supplied by the peripheral storage device or systems vendor.

4.2. OPERATING SYSTEM CONSIDERATIONS

4.2.1. Windows Operating System Considerations for ATA Drives

4.2.1.1. Windows 3.1

Under Microsoft Windows 3.1 two possible configurations may be present. In the first, 32-bit Disk Access (also known as Fastdisk) may be disabled. In this case access to the task file registers with the operation of the 32-bit Disk Access is disabled. This prevents access to the task file registers. The actual device changes to the 32-bit Disk Access (Fastdisk) ROM or some NTFS or FAT32 BIOS mechanism for S.M.A.R.T. commands. Use of an additional 32-bit virtual device driver (VxD) would also provide a method to access the task file registers. As with the case where 32-bit Disk Access is disabled, care must be taken to synchronize access to the task file.

4.2.2.1. Windows 95 and Windows NT

Programmers writing Win32 applications, for the Windows 95 and Windows NT platforms, should use the Compaq/Microsoft S.M.A.R.T. IOCTL API Specification referenced in Section 2.

5. HARD DISK DRIVE IMPLEMENTATION GUIDELINES

5.1. ATA HARD DISK DRIVE GUIDELINES

Refer to the ATA-3 specification for information about other features of the S.M.A.R.T. system. In order for the software to be compatible across all vendors and all disk drives, developers are strongly encouraged to limit their software to the functions listed in Table 1 of this document. Other functions in the ATA-3 document encompass vendor specific elements that may cause inconsistent results.

Note that S.M.A.R.T. sub-commands D5h through D6h and DBh through DFh are reserved, and sub-commands D7h and E0h through EFh are vendor specific.

5.2. SCSI HARD DISK DRIVE GUIDELINES

Applications or BIOS may wish to enable S.M.A.R.T., assess the S.M.A.R.T. status and then pass control on to another application or driver. Drives must report accurate S.M.A.R.T. status for the first successful command after S.M.A.R.T. is enabled.

SCSI S.M.A.R.T. drives may autosave S.M.A.R.T. data to disk when the Perf bit in Mode Page 1Ch is set to 0.

6. SOFTWARE APPLICATIONS GUIDE COMPLIANCE

Applications, BIOS and drivers that are compliant with this specification may display the following S.M.A.R.T. compliant logo on their product and associated literature, without payment of royalties or legal agreements.

ATA Disk devices that implement the S.M.A.R.T. sub-commands and their respective sub-command codes listed in Table 1 may display the S.M.A.R.T. compliant logo.

SCSI Disk devices that implement the Informational Exception mode page and the reporting method listed in Table 2 may display the S.M.A.R.T. compliant logo, (see Section 2., SCSI Physical Commands reference).



S.M.A.R.T. information, including this logo, is available from member companies.

APPENDIX

EXECUTE OFF-LINE IMMEDIATE (Optional)

Command Description

INPUTS - The Features register shall be set to D4h. The Cylinder Low register shall be set to 4Fh. The Cylinder High register shall be set to C2h.

NORMAL OUTPUTS - None

ERROR OUTPUTS - If the device does not support this subcommand, if S.M.A.R.T. is disabled or if the values in the Features, Cylinder Low or Cylinder High registers are invalid, an Aborted command error is posted.

PREREQUISITES - DRDY set equal to one. S.M.A.R.T. enabled.

DESCRIPTION - This subcommand causes the device to immediately initiate or resume the optional set of activities that collect Attribute data in an off-line mode and then save this data to the device's non-volatile memory.

If the S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand is supported by the device: upon receipt of the subcommand from the host, the device sets BSY to one, begins or resumes its set of off-line activities, clears BSY to zero and asserts INTRQ.

During execution of its off-line activities the device shall not set BSY nor clear DRDY.

If the device is in the process of performing its set of off-line activities as a result of receiving a S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand from the host and is interrupted by any new command from the host except a S.M.A.R.T. DISABLE OPERATIONS, S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE or STANDBY IMMEDIATE command, the device shall suspend or abort its off-line activities and service the host within two seconds after receipt of the new command. After servicing the interrupting command from the host the device may immediately re-initiate or resume off-line activities without any additional commands from the host.

If the device is in the process of performing its off-line activities and is interrupted by a STANDBY IMMEDIATE command from the host, the device shall suspend or abort its off-line activities, and service the host within two seconds after receipt of the command. After receiving a new command that causes the device to exit a power saving mode, the device shall either wait for a new S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand from the host to re-initiate or resume off-line activities, or the device shall immediately re-initiate or resume off-line activities without any additional commands from the host.

If the device is in the process of performing its off-line activities and is interrupted by a S.M.A.R.T. DISABLE OPERATIONS subcommand from the host, the device shall suspend or abort its off-line activities and service the host within two seconds after receipt of the command. Upon receipt of the next S.M.A.R.T. ENABLE OPERATIONS subcommand the device may, after the next vendor specified event, either re-initiate its off-line activities or resume those activities from where they had been previously suspended.

If the device is in the process of performing its off-line activities and is interrupted by a S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand from the host, the device shall abort its off-line activities and service the host within two seconds after receipt of the command. The device will then re-initiate its off-line data collection activity in response to the new EXECUTE OFF-LINE IMMEDIATE subcommand.

Status and Information Bytes

For the EXECUTE OFF-LINE IMMEDIATE subcommand to be effectively used, application software needs to be aware of certain capabilities of the device regarding off-line data collection. Several bytes of information that apply to the EXECUTE OFF-LINE IMMEDIATE subcommand are returned in the SMART READ ATTRIBUTE VALUES data buffer. These status and information bytes are listed below, along with their definitions.

Table 2 -- OFF-LINE IMMEDIATE Status and Information Bytes

Description	Byte s	Byte Offset	Type
Off-Line Data Collection Status	1	362	Rd only
Time in seconds to complete Off-Line data collection activity	2	364	Rd only
Off-Line data collection capability	1	367	Rd only

The value of the off-line data collection status byte defines the current status of the off-line activities of the device. The values and their respective definitions are listed in Table 3.

Table 3 - Off-line data collection status byte values

Value	Definition
00h	Off-line data collection activity was never started.
01h	Reserved
02h	Off-line data collection activity was completed without error.
03h	Reserved
04h	Off-line data collection activity was suspended by an interrupting command from host.
05h	Off-line data collection activity was aborted by an interrupting command from host.
06h	Off-line data collection activity was aborted by the device with a fatal error.
07h thru 3Fh	Reserved
40h thru	Vendor specific

7Fh	
80h	Off-line data collection activity was never started. The auto-off-line feature is enabled.
81h	Reserved
82h	Off-line data collection activity was completed without error. The auto-off-line feature is enabled.
83h	Reserved
84h	Off-line data collection activity was suspended by an interrupting command from host. The auto-off-line feature is enabled.
85h	Off-line data collection activity was aborted by an interrupting command from host. The auto-off-line feature is enabled.
86h	Off-line data collection activity was aborted by the device with a fatal error. The auto-off-line feature is enabled.
87h thru BFh	Reserved
C0h thru FFh	Vendor specific

Time in seconds to complete off-line data collection activity

This word specifies how many seconds the device requires to complete its sequence of off-line data collection activity. Valid values for this word are from 0001h to FFFFh.

Off-line data collection capability

The following describes the definition for the off-line data collection capability bits. If the value of all of these bits is equal to zero, then no off-line data collection is implemented by this device.

- Bit 0 (EXECUTE OFF-LINE IMMEDIATE implemented bit) - If the value of this bit equals one, then the S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand is implemented by this device. If the value of this bit equals zero, then the S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand is not implemented by this device.
- Bit 1 (ENABLE/DISABLE AUTOMATIC OFF-LINE implemented bit) - If the value of this bit equals one, then the S.M.A.R.T. ENABLE/DISABLE AUTOMATIC OFF-LINE subcommand is implemented by this device. If the value of this bit equals zero, then the S.M.A.R.T. ENABLE/DISABLE AUTOMATIC OFF-LINE subcommand is not implemented by this device.
- Bit 2 (abort/restart off-line by host bit) - If the value of this bit equals one, then the device shall abort all off-line data collection activity initiated by an S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE command upon receipt of a new command. Off-line data collection activity must be restarted by a new S.M.A.R.T. EXECUTE OFF-LINE IMMEDIATE subcommand from the host. If the value of this bit equals zero, the device shall suspend off-line data collection activity after an interrupting command and resume off-line data collection activity after some vendor-specified event.
- Bits 3-7 (reserved bits) - All bits other not defined in this section are reserved for future use.

Implementation of this feature is optional and vendor specific.