Vysoké učení technické v Brně
Fakulta informačních technologií
ITP – Technika personálních počítačů
Analýza rozhraní PATA a SATA, S.M.A.R.T.

Úloha č.:
10

Zadání:

1. Seznamte se s ovládáním programu BusTrace.
2. Prostřednictvím nástroje „Správa disků“ ve Windows na pokusném disku vytvořte logický oddíl. Komunikaci při vytváření oddílu zaznamenejte.
3. Vytvořený oddíl naformátujte na FAT32. Použijte rychlé formátování a generovanou komunikaci zaznamenejte.
    a. Zapisuje formátování data na začátek disku?
4. Na pokusný disk nakopírujte textový soubor. Analýzou komunikace zjistěte, na které umístění se soubor nahrál. Zjisštěnou informaci ověřte pomocí programu WinHex.
5. Prostudujte technologii S.M.A.R.T. pro předpovídání selhání disků.
6. Seznamte se s ovládáním programu smart tools. Program se spouští z příkazové řádky příkazem smartctl.
7. S pomocí programu smart tools zjistěte stav disku.
    a. Kolik hodin byl již disk v provozu?
    b. Je tato hodnota spolehlivá?
8. Spusťte krátký test disku a zaznamenejte komunikaci s diskem.
9. Zaznamenanou komunikaci důkladně analyzujte a vysvětlete význam jednotlivých příkazů.
10. Zrušte pokusný oddíl a vypněte počítač.

S.M.A.R.T.

Self-Monitoring, Analysis, and Reporting Technology (**S.M.A.R.T.**) je v informatice monitorovací systém pro pevné disky. Firmware sleduje uvnitř pevného disku během provozu různé parametry, které je možné přenést do počítače. Pomocí zjištěných údajů může správce systému včas předvídat selhání disku a uložená data zálohovat. Systém S.M.A.R.T. definuje hlavně způsob komunikace disku s počítačem, takže jednotlivé implementace se u různých výrobců liší.

Poruchy pevných disků

Poruchy pevných disků můžeme rozdělit do dvou hlavních kategorií:

Předvídatelné poruchy

> Předvídatelné poruchy většinou gradují postupem času (například mechanické opotřebení). Monitorovací zařízení umí zjistit tyto problémy a signalizovat jejich stav, podobně jako kontrolka na palubní desce auta signalizuje přehřátí motoru. Toto monitorování může pomoci předejít daným problémům, nebo umožnit jejich nápravu dříve, než bude příliš pozdě.

Nepředvídatelné poruchy

> Nepředvídatelné poruchy nastávají náhle, bez předchozích náznaků. Důvod může být různý, od poškození elektroniky disku, po fyzickou destrukci (např. vlivem silných otřesů - pád disku a podobně).

Mechanické poškození je zodpovědné za 60 % všech selhání.[1] Většině těchto poruch, ale předchází množství indikátorů, jako například zvýšené zahřívání disku, zvýšený hluk, problémy s čtením či zápisem dat, zvyšující se počet poškozených sektorů, a podobně. Při opakovaném monitorování těchto příznaků lze tedy s nezanedbatelnou pravděpodobností určit, zda vývoj hodnot značí blízkou poruchu.

Zaměstnanci firmy Google ve své práci zjistili několik parametrů disku, které mají skutečně úzký vztah k poruchám disků, nicméně dle jejich měření odhady, založené čistě na S.M.A.R.T. atributech jsou zřídkakdy užitečné pro předpověď selhání individuálního disku. Při nasazení na rozsáhlejší populaci disků je jejich vypovídací hodnota vyšší, ale pro vytvoření spolehlivého modelu předpovědí nejsou tyto informace samy o sobě dostačující. Tvůrci práce dále poukazují například na to, že vyšší teplota, nebo časté užívání nemají zřejmě tak velký vliv, jak bylo dříve naznačováno.[2]

Příklady poruch a jejich indikátorů:

Typ poruchy	Indikace
poškození povrchu disku	vzrůstající počet vadných/přemapovaných sektorů
poškození hlavy	nárůst měkkých chyb („soft errors“), opakované pokusy o čtení, ECC chyby
poškození motoru	vibrace, zvýšený hluk motoru, ložisek, nárůst času potřebného k roztočení disku
poškození elektroniky disku	žádné (náhlá porucha)
poškození vystavovacího mechanismu	chyby při vyhledávání uložených dat („seek“)

Historie

První monitorovací technologie byla představena IBM v roce 1992 v jejich diskových polích IBM 9337 používajících SCSI-2 disky. Tato technologie byla pojmenována *Predictive Failure Analysis* (PFA, analýza předvídatelných poruch). Metoda spočívala v měření několika parametrů, klíčových pro spolehlivost disku, a jejich vyhodnocování ve firmwaru disku. Komunikace mezi fyzickým diskem a monitorovacím softwarem byla omezená pouze na 2 stavy: disk je v pořádku nebo disk pravděpodobně brzo selže.

Později byla firmami Compaq, Seagate, Quantum a Conner vytvořena technologie IntelliSafe. Sledovány byly parametry „zdraví“ (spolehlivosti) disku, tyto hodnoty byly přenášeny do operačního systému a zobrazovány monitorovacím softwarem. Každý výrobce disků se mohl sám rozhodnout, jaké parametry zahrne do monitorování, i jaké pro ně budou prahové hodnoty. Sjednocení bylo na úrovni protokolu komunikace mezi diskem a počítačem.

Compaq předložil svou implementaci na počátku roku 1995 komisi pro standardizaci, která technologii IntelliSafe schválila. Výsledný standard byl pojmenován S.M.A.R.T.

S.M.A.R.T.

Nejzákladnější informace, kterou S.M.A.R.T. poskytuje, je S.M.A.R.T. status. Status nabývá pouze dvou hodnot: „prahová hodnota nepřekročena“ a „prahová hodnota překročena“. Tyto stavy jsou také často reprezentovány jako „disk je v pořádku“ a „disk selhal“ (myšleno z hlediska S.M.A.R.T. testu), respektive stav „prahová hodnota překročena“ lze reprezentovat jako relativně vysokou šanci, že disk v blízké budoucnosti selže. Předpovídané selhání může být fatální porucha, nebo také pouhé snížení výkonnosti kvůli drobným problémům.

Detailnější pohled na stav disku je možné získat prozkoumáním S.M.A.R.T. Atributů. S.M.A.R.T. Atributy byly obsaženy v některých návrzích ATA standardů, ale byly odstraněny z finálního standardu. Monitorování a interpretace S.M.A.R.T. Atributů se totiž u jednotlivých výrobců liší. Nelze tedy přímo podle atributů porovnávat disky různých výrobců, jelikož disky nemusí monitorovat stejné atributy, a nebo mají nastavené jiné interpretace a prahové hodnoty stejných atributů. Disky s podporou S.M.A.R.T. technologie mohou také také umožňovat „logování“ událostí a chyb spojených s diskem. Tyto záznamy („logy“) mohou pomoci při údržbě a servisu počítače (např. určit, zda příčinou problémů s počítačem je pevný disk).

Další vlastností některých disků podporujících S.M.A.R.T. technologii, je možnost testovat daný disk (tzv. On-Self-Test), například na zjištění chyb povrchu disku (vadných sektorů). Výsledky těchto testů se také většinou ukládají do „logů“ a jejich analýzou lze zjistit stav disku.

Všechny tyto informace S.M.A.R.T. technologie slouží k minimalizování rizika ztráty dat. Smysl těchto informací spočívá v predikci selhání a umožnění zálohy dat dříve, než k němu dojde. S.M.A.R.T. technologie je technologií informační, neobsahuje prostředky k obnově dat v případě, že k selhání již došlo. To znamená, že S.M.A.R.T. technologie není nic platná v případě náhlého selhání, v případě že informace že informace nejsou uživatelem správně vyhodnoceny a v některých případech nemusí S.M.A.R.T. indikovat správný stav disku (disk označený jako nespolehlivý může fungovat bez vážných problémů ještě několik let a naopak).

Mezi výrobce podporující S.M.A.R.T. patří: Samsung, Seagate, IBM (Hitachi), Fujitsu, Maxtor, Toshiba, Western Digital and ExcelStor Technology.

Atributy

Každý výrobce disků si definuje vlastní sadu atributů a stanoví jejich prahové hodnoty, k jejichž překročení by nemělo za normálních podmínek dojít (práh může nabývat hodnot 0 - 253, při nule nemůže atribut selhat - nezapočítává se do S.M.A.R.T. statusu). Atributy mají takzvanou “syrovou hodnotu” (raw value) která často označuje fyzickou hodnotu dané veličiny (např. stupně v Celsia,počet vadných sektorů,sekundy) a normalizovanou hodnotu (v rozmezí 1/nejhorší - 253/nejlepší) odvozenou podle aktuálního stavu (syrové hodnoty), která je porovnávána s hodnotou,určenou jako normální (tzn. bez jakýchkoli problémů). Nastavení těchto hodnot a prahu je ovšem čistě věcí výrobce, tudíž dva různé disky mohou zobrazovat při stejném fyzickém stavu jiné hodnoty pro daný atribut.

- Příklad 1:

ID Atributu je 4 („počet roztočení disku“): Hodnota je 253, nejhorší zaznamenaná hodnota je 253, práh je 0, syrová hodnota je 1324. Jelikož nedochází ke zhoršování hodnoty (hodnota je rovna nejhorší naměřené) a práh je 0, je tento atribut pouze informační a nemá vliv na S.M.A.R.T. status disku. Syrová hodnota v tomto případě udává počet

roztočení disku.

- Příklad 2:

ID Atributu je 5 („počet přemapovaných sektorů“): Hodnota je 112, nejhorší zaznamenaná hodnota je 112, práh je 63, syrová hodnota je 3. V tomto případě má již atribut přímou souvislost se spolehlivostí disku (práh má hodnotu 63), ale hodnota 112 je nad prahovou hodnotou, takže lze předpokládat že disku nehrozí v blízké době žádné selhání, ale není již v bezvadném stavu (hodnota atributu je pod 235, syrová hodnota značí 3 přemapované sektory). Pokud by se počet přemapovaných sektorů zvětšoval, mohlo by to pravděpodobně znamenat blížící se selhání disku a bylo by tedy vhodné provést zálohu.

Zde je uveden a přeložen report nejběžnějších chybových hlášení tabulky S.M.A.R.T.

Reference

[1] Seagate Product Marketing. *Get S.M.A.R.T. for Reliability* [online]. California, USA : Seagate, rev. 1999-07-13, [cit. 2009-09-06]. Dostupné online. (http://www.seagate.com/docs/pdf/whitepaper/enhanced_smart.pdf) (anglicky)

[2] Failure Trends in a Large Disk Drive Population (http://labs.google.com/papers/disk_failures.pdf) (Conclusion section), by Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz André Barroso, Google Inc. 1600 Amphitheatre Pkwy Mountain View, CA 94043

S.M.A.R.T.

From Wikipedia, the free encyclopedia

S.M.A.R.T. (**Self-Monitoring, Analysis, and Reporting Technology**; sometimes written as **SMART**) is a monitoring system for computer hard disk drives to detect and report on various indicators of reliability, in the hope of anticipating failures.

When a failure is anticipated by S.M.A.R.T., the drive should be replaced, and can sometimes be returned to the manufacturer, who can use these failed drives to discover where faults lie and how to prevent them from recurring on the next generation of hard disk drives.

Contents

Background

The purpose of S.M.A.R.T. is to warn a user or a system administrator of impending drive failure while there is still time to take action, such as copying the data to a replacement device.

Hard disk failures fall into one of two basic classes:

- Predictable failures: These failures result from slow processes such as mechanical wear and gradual degradation of storage surfaces. Monitoring can determine when such failures are becoming more likely.
- Unpredictable failures: These failures happen suddenly and without warning. They range from electronic components becoming defective to a sudden mechanical failure (perhaps due to improper handling).

Mechanical failures account for about 60% of all drive failures.[1] While the eventual failure may be catastrophic, most mechanical failures result from gradual wear and there are usually certain indications that failure is imminent. These may include increased heat output, increased noise level, problems with reading and writing of data, an increase in the number of damaged disk sectors.

Work at Google on over 100,000 drives has shown little overall predictive value of S.M.A.R.T. status as a whole, but suggests that certain sub-categories of information which some S.M.A.R.T. implementations track *do* correlate with actual failure rates: specifically, in the 60 days following the first scan error on a drive, the drive is, on average, 39 times more likely to fail than it would have been had no such error occurred. Furthermore, first errors in reallocations, offline reallocations and probational counts are strongly correlated to higher probabilities of failure.[2]

PCTechGuide's page on S.M.A.R.T. (2003)[3] comments that the technology has gone through three phases:

> "In its original incarnation SMART provided failure prediction by monitoring certain online hard drive activities. A subsequent version improved failure prediction by adding an automatic off-line read scan to monitor additional operations. The latest "SMART" technology not only monitors hard drive activities but adds failure prevention by attempting to detect and repair sector errors. Also, while earlier versions of the technology only monitored hard drive activity for data that was retrieved by the operating system, this latest SMART tests all data and all sectors of a drive by using "off-line data collection" to confirm the drive's health during periods of inactivity."

History and predecessors

An early hard disk monitoring technology was introduced by IBM in 1992 in their IBM 9337 Disk Arrays for AS/400 servers using IBM 0662 SCSI-2 disk drives.[4] Later it was named Predictive Failure Analysis (PFA) technology. It was measuring several key device health parameters and evaluating them within the drive firmware. Communications between the physical unit and the monitoring software were limited to a binary result: namely, either "device is OK" or "drive is likely to fail soon".

Later, another variant, which was named IntelliSafe, was created by computer manufacturer Compaq and disk drive manufacturers Seagate, Quantum, and Conner.[5] The disk drives would measure the disk's "health parameters", and the values would be transferred to the operating system and user-space monitoring software. Each disk drive vendor was free to decide which parameters were to be included for monitoring, and what their thresholds should be. The unification was at the protocol level with the host.

Compaq submitted their implementation to Small Form Committee for standardization in early 1995.[6] It was supported by IBM, by Compaq's development partners Seagate, Quantum, and Conner, and by Western Digital, who did not have a failure prediction system at the time. The Committee chose IntelliSafe's approach, as it provided more flexibility. The resulting jointly developed standard was named S.M.A.R.T.

Information

The technical documentation for SMART is in the AT Attachment (ATA) standard.[7]

The most basic information that SMART provides is the SMART status. It provides only two values: "threshold not exceeded" and "threshold exceeded". Often these are represented as "drive OK" or "drive fail" respectively. A "threshold exceeded" value is intended to indicate that there is a relatively high probability that the drive will not be able to honor its specification in the future: that is, the drive is "about to fail". The predicted failure may be catastrophic or may be something as subtle as the inability to write to certain sectors, or perhaps slower performance than the manufacturer's declared minimum.

The SMART status does not necessarily indicate the drive's past or present reliability. If a drive has already failed catastrophically, the SMART status may be inaccessible. Alternatively, if a drive has experienced problems in the past, but the sensors no longer detect such problems, the SMART status may, depending on the manufacturer's programming, suggest that the drive is now sound.

The inability to *read* some sectors is not always an indication that a drive is about to fail. One way that unreadable sectors may be created, even when the drive is functioning within specification, is through a sudden power failure while the drive is writing. In order to prevent this problem, modern hard drives will always finish writing at least the current sector immediately after the power fails (typically using rotational energy from the disk). Also, even if the physical disk is damaged at one location, such that a certain sector is unreadable, the disk may be able to use spare space to replace the bad area, so that the sector can be overwritten.[8]

More detail on the health of the drive may be obtained by examining the SMART Attributes. SMART Attributes were included in some drafts of the ATA standard, but were removed before the standard became final. The meaning and interpretation of the attributes varies between manufacturers, and are sometimes considered a trade secret for one manufacturer or another. Attributes are further discussed below.[9]

Drives with SMART may optionally support a number of 'logs'. The *error log* records information about the most recent errors that the drive has reported back to the host computer. Examining this log may help one to determine whether computer problems are disk-related or caused by something else.

A drive supporting SMART may optionally support a number of self-test or maintenance routines, and the results of the tests are kept in the *self-test log*. The self-test routines may be used to detect any unreadable sectors on the disk, so that they may be restored from back-up sources (for example, from other disks in a RAID). This helps to reduce the risk of incurring permanent loss of data.

Standards and implementation

Many motherboards will display a warning message when a disk drive is approaching failure. Although an industry standard exists among most major hard drive manufacturers,[3] there are some remaining issues and much proprietary "secret knowledge" held by individual manufacturers as to their specific approach. As a result, S.M.A.R.T. is not always implemented correctly on many computer platforms, due to the absence of industry-wide software and hardware standards for S.M.A.R.T. data interchange.[*citation needed*]

From a legal perspective, the term "S.M.A.R.T." refers only to a signaling method between internal disk drive electromechanical sensors and the host computer. Hence, a drive may be claimed by its manufacturers to include S.M.A.R.T. support even if it does not include, say, a temperature sensor, which the customer might reasonably expect to be present. Moreover, in the most extreme case, a disk manufacturer could, in theory, produce a drive which includes a sensor for just *one* physical attribute, and then legally advertise the product as "S.M.A.R.T. compatible".

Depending on the type of interface being used, some S.M.A.R.T.-enabled motherboards and related software may not communicate with certain S.M.A.R.T.-capable drives. For example, few external drives connected via USB and Firewire correctly send S.M.A.R.T. data over those interfaces. With so many ways to connect a hard drive (SCSI, Fiber Channel, ATA, SATA, SAS, SSA, and so on), it is difficult to predict whether S.M.A.R.T. reports will function correctly in a given system.

Even on hard drives and interfaces that support it, S.M.A.R.T. information may not be reported correctly to the computer's operating system. Some disk controllers can duplicate all write operations on a secondary "back-up" drive in real time. This feature is known as "RAID mirroring". However, many programs which are designed to analyze changes in drive behavior and relay S.M.A.R.T. alerts to the operator do not function properly when a computer system is configured for RAID support. Generally this is because, under normal RAID operational conditions, the computer is not permitted by the RAID subsystem to 'see' (or directly access) individual physical drives, but may access only logical volumes instead.

On the Windows platform, many programs designed to monitor and report S.M.A.R.T. information will function only under an administrator account. At present, S.M.A.R.T. is implemented individually by manufacturers, and while some aspects are standardized for compatibility, others are not.

Access

For a list of various programs that allow reading of Smart Data, see Comparison of S.M.A.R.T. tools.

ATA S.M.A.R.T. attributes

Each drive manufacturer defines a set of attributes, and sets threshold values beyond which attributes should not pass under normal operation. Each attribute has a *raw value*, whose meaning is entirely up to the drive manufacturer (but often corresponds to counts or a physical unit, such as degrees Celsius or seconds), a normalized value, which ranges from 1 to 253 (with 1 representing the worst case and 253 representing the best) and a worst value, which represents the lowest recorded normalized value. Depending on the manufacturer, a value of 100 or 200 will often be chosen as the "normal" value.

Manufacturers that have supported at least one S.M.A.R.T. attribute in various products include: Samsung, Seagate, IBM (Hitachi), Fujitsu, Maxtor, Toshiba, Intel, Western Digital and ExcelStor Technology.

Known ATA S.M.A.R.T. attributes

The following chart lists some S.M.A.R.T. attributes and the typical meaning of their raw values. Normalized values are always mapped so that higher values are better (with only very rare exceptions such as the "Temperature" attribute on certain Seagate drives[10]), but higher *raw* attribute values may be better or worse depending on the attribute and manufacturer. For example, the "Reallocated Sectors Count" attribute's normalized value *decreases* as the count of reallocated sectors *increases*. In this case, the attribute's *raw* value will often indicate the actual count of sectors that were reallocated, although vendors are in no way required to adhere to this convention. As manufacturers do not necessarily agree on precise attribute definitions and measurement units, the following list of attributes should be regarded as a general guide only.

Legend

	Higher raw value is better
	Lower raw value is better
Critical: red colored row	Potential indicators of imminent electromechanical failure

ID	Hex	Attribute name	Better	Description
01	0x01	**Read Error Rate**		Stores data related to the rate of hardware read errors that occurred when reading data from a disk surface. The raw value has different structure for different vendors and is often not meaningful as a decimal number.
02	0x02	**Throughput Performance**		Overall (general) throughput performance of a hard disk drive. If the value of this attribute is decreasing there is a high probability that there is a problem with the disk.
03	0x03	**Spin-Up Time**		Average time of spindle spin up (from zero RPM to fully operational [millisecs]).
04	0x04	**Start/Stop Count**		A tally of spindle start/stop cycles. The spindle turns on, and hence the count is increased, both when the hard disk is turned on after having before been turned entirely off (disconnected from power source) and when the hard disk returns from having previously been put to sleep mode.[11]
05	0x05	**Reallocated Sectors Count**		Count of reallocated sectors. When the hard drive finds a read/write/verification error, it marks this sector as "reallocated" and transfers data to a special reserved area (spare area). This process is also known as remapping, and "reallocated" sectors are called remaps. Unfortunately, on modern operating systems, such as Windows XP and onwards, "bad blocks" cannot be found while testing the surface, as this feature was removed. However, 3rd-party applications such as "HD Tune" can reveal bad sectors across the entire surface, even on partitions that are hidden. Also, as the count of reallocated sectors increases, the read/write speed tends to decrease, unless the bad sectors are manually repositioned to a hidden partition, although the boot sector is always at the start of the disk, so if damage is in that area, the drive is only useful as a redundant backup drive. The raw value normally represents a count of the count of bad sectors that have been found and remapped. Thus, the higher the attribute value, the more sectors the drive has had to reallocate.
06	0x06	**Read Channel Margin**		Margin of a channel while reading data. The function of this attribute is not specified.
07	0x07	**Seek Error Rate**	N/A	(Vendor specific raw value.) Rate of seek errors of the magnetic heads. If there is a partial failure in the mechanical positioning system, then seek errors will arise. Such a failure may be due to numerous factors, such as damage to a servo, or thermal widening of the hard disk. The raw value has different structure for different vendors and is often not meaningful as a decimal number.
08	0x08	**Seek Time Performance**		Average performance of seek operations of the magnetic heads. If this attribute is decreasing, it is a sign of problems in the mechanical subsystem.
09	0x09	**Power-On Hours (POH)**		Count of hours in power-on state. The raw value of this attribute shows total count of hours (or minutes, or seconds, depending on manufacturer) in power-on state.[12]
10	0x0A	**Spin Retry Count**		Count of retry of spin start attempts. This attribute stores a total count of the spin start attempts to reach the fully operational speed (under the condition that the first attempt was unsuccessful). An increase of this attribute value is a sign of problems in the hard disk mechanical subsystem.
11	0x0B	**Recalibration Retries Calibration_Retry_Count**		This attribute indicates the count that recalibration was requested (under the condition that the first attempt was unsuccessful). An increase of this attribute value is a sign of problems in the hard disk mechanical subsystem.
12	0x0C	**Power Cycle Count**		This attribute indicates the count of full hard disk power on/off cycles.
13	0x0D	**Soft Read Error Rate**		Uncorrected read errors reported to the operating system.
183	0xB7	**SATA Downshift Error Count**		Western Digital and Samsung attribute.
184	0xB8	**End-to-End error**		This attribute is a part of HP's SMART IV technology and it means that after transferring through the cache RAM data buffer the parity data between the host and the hard drive did not match.[13]
185	0xB9	**Head Stability**		Western Digital attribute.
186	0xBA	**Induced Op-Vibration Detection**		Western Digital attribute.
187	0xBB	**Reported Uncorrectable Errors**		The count of errors that could not be recovered using hardware ECC (see attribute 195).
188	0xBC	**Command Timeout**		The count of aborted operations due to HDD timeout. Normally this attribute value should be equal to zero and if the value is far above zero, then most likely there will be some serious problems with power supply or an oxidized data cable.[13]
189	0xBD	**High Fly Writes**		HDD producers implement a Fly Height Monitor that attempts to provide additional protections for write operations by detecting when a recording head is flying outside its normal operating range. If an unsafe fly height condition is encountered, the write process is stopped, and the information is rewritten or reallocated to a safe region of the hard drive. This attribute indicates the count of these errors detected over the lifetime of the drive. This feature is implemented in most modern Seagate drives[1] and some of Western Digital's drives, beginning with the WD Enterprise WDE18300 and WDE9180 Ultra2 SCSI hard drives, and will be included on all future WD Enterprise products.[14]
190	0xBE	**Airflow Temperature (WDC)**		Airflow temperature on Western Digital HDs (Same as temp. [C2], but current value is 50 less for some models. Marked as obsolete.)
190	0xBE	**Temperature Difference from 100**		Value is equal to (100−temp. °C), allowing manufacturer to set a minimum threshold which corresponds to a maximum temperature. (Seagate only?)[*citation needed*] Seagate ST910021AS: Verified Present[*citation needed*] Seagate ST9120823ASG: Verified Present under name "Airflow Temperature Cel" 2008-10-06 Seagate ST3802110A: Verified Present 2007-02-13[*citation needed*]

				Seagate ST980825AS: Verified Present 2007-04-05[*citation needed*]
				Seagate ST3320620AS: Verified Present 2007-04-23[*citation needed*]
				Seagate ST3500641AS: Verified Present 2007-06-12[*citation needed*]
				Seagate ST3250824AS: Verified Present 2007-08-07[*citation needed*]
				Seagate ST3250620AS: Verified Present
				Seagate ST31000340AS: Verified Present 2008-02-05[*citation needed*]
				Seagate ST31000333AS: Verified Present 2008-11-24[*citation needed*]
				Seagate ST3160211AS: Verified Present 2008-06-12[*citation needed*]
				Seagate ST3320620AS: Verified Present 2008-06-12[*citation needed*]
				Seagate ST3400620AS: Verified Present 2008-06-12[*citation needed*]
				Seagate ST3750330AS: Verified present 2009-07-06[*citation needed*]
				Seagate ST3500418AS: Verified present 2010-04-03
				Seagate ST31500341AS: Verified present 2010-10-09
				Samsung HD501LJ: Verified Present under name "Airflow Temperature" 2008-03-02[*citation needed*]
				Samsung HD753LJ: Verified Present under name "Airflow Temperature" 2008-07-15[*citation needed*]
				
				A note here: smartctl seems to interpret these correctly at least in 5.39.1:
				
				`ID# ATTRIBUTE_NAME FLAG VALUE WORST THRESH TYPE UPDATED WHEN_FAILED RAW_VALUE`
				`190 Airflow_Temperature_Cel 0x0022 068 057 045 Old_age Always - 32 (Lifetime Min/Max 22/33)`
				
				notice "raw value" is 32 (the correct airflow temp in celsius) and value is 100-32 = 68.

191	0xBF	**G-sense error rate**		The count of errors resulting from externally-induced shock & vibration.
192	0xC0	**Power-off Retract Count** **Emergency Retract Cycle count** (Fujitsu)[15]		Count of times the heads are loaded off the media. Heads can be unloaded without actually powering off.[*citation needed*]
193	0xC1	**Load Cycle Count** **Load/Unload Cycle Count** (Fujitsu)		Count of load/unload cycles into head landing zone position.[15] The typical lifetime rating for laptop (2.5-in) hard drives is 300,000 to 600,000 load cycles.[16] Some laptop drives are programmed to unload the heads whenever there has not been any activity for about five seconds.[17] Many Linux installations write to the file system a few times a minute in the background.[18] As a result, there may be 100 or more load cycles per hour, and the load cycle rating may be exceeded in less than a year.[19]
194	0xC2	**Temperature**		Current internal temperature.
195	0xC3	**Hardware ECC Recovered**	N/A	(Vendor specific raw value.) The raw value has different structure for different vendors and is often not meaningful as a decimal number.
196	0xC4	**Reallocation Event Count**		Count of remap operations. The raw value of this attribute shows the total count of attempts to transfer data from reallocated sectors to a spare area. Both successful & unsuccessful attempts are counted.[20]
197	0xC5	**Current Pending Sector Count**		Count of "unstable" sectors (waiting to be remapped, because of read errors). If an unstable sector is subsequently written or read successfully, this value is decreased and the sector is not remapped. Read errors on a sector will not remap the sector (since it might be readable later); instead, the drive firmware remembers that the sector needs to be remapped, and remaps it the next time it's written.[21]
198	0xC6	**Uncorrectable Sector Count**		The total count of uncorrectable errors when reading/writing a sector. A rise in the value of this attribute indicates defects of the disk surface and/or problems in the mechanical subsystem. (or Off-Line Scan Uncorrectable Sector Count: Fujitsu)[15]
199	0xC7	**UltraDMA CRC Error Count**		The count of errors in data transfer via the interface cable as determined by ICRC (Interface Cyclic Redundancy Check).
200	0xC8	**Multi-Zone Error Rate** [22]		The count of errors found when writing a sector. The higher the value, the worse the disk's mechanical condition is.
200	0xC8	**Write Error Rate** (Fujitsu)		The total count of errors when writing a sector.[23]
201	0xC9	**Soft Read Error Rate**		Count of off-track errors.
202	0xCA	**Data Address Mark errors**		Count of Data Address Mark errors (or vendor-specific).[*citation needed*]
203	0xCB	**Run Out Cancel**		Count of ECC errors
204	0xCC	**Soft ECC Correction**		Count of errors corrected by software ECC[*citation needed*]
205	0xCD	**Thermal Asperity Rate (TAR)**		Count of errors due to high temperature.[13]
206	0xCE	**Flying Height**		Height of heads above the disk surface. A flying height that's too low increases the chances of a head crash while a flying height that's too high increases the chances of a read/write error.[*citation needed*]
207	0xCF	**Spin High Current**		Amount of surge current used to spin up the drive.[13]

208	0xD0	**Spin Buzz**		Count of buzz routines needed to spin up the drive due to insufficient power.[13]
209	0xD1	**Offline Seek Performance**		Drive's seek performance during its internal tests.[13]
210	0xD2	?		(found in a Maxtor 6B200M0 - 200GB disk)
211	0xD3	**Vibration During Write**		Vibration During Write[*citation needed*]
212	0xD4	**Shock During Write**		Shock During Write[*citation needed*]
220	0xDC	**Disk Shift**		Distance the disk has shifted relative to the spindle (usually due to shock or temperature). Unit of measure is unknown.
221	0xDD	**G-Sense Error Rate**		The count of errors resulting from externally-induced shock & vibration.
222	0xDE	**Loaded Hours**		Time spent operating under data load (movement of magnetic head armature)[*citation needed*]
223	0xDF	**Load/Unload Retry Count**		Count of times head changes position.[*citation needed*]
224	0xE0	**Load Friction**		Resistance caused by friction in mechanical parts while operating.[*citation needed*]
225	0xE1	**Load/Unload Cycle Count**		Total count of load cycles[*citation needed*]
226	0xE2	**Load 'In'-time**		Total time of loading on the magnetic heads actuator (time not spent in parking area).[*citation needed*]
227	0xE3	**Torque Amplification Count**		Count of attempts to compensate for platter speed variations[*citation needed*]
228	0xE4	**Power-Off Retract Cycle**		The count of times the magnetic armature was retracted automatically as a result of cutting power.[*citation needed*]
230	0xE6	**GMR Head Amplitude**		Amplitude of "thrashing" (distance of repetitive forward/reverse head motion)[*citation needed*]
231	0xE7	**Temperature**		Drive Temperature
240	0xF0	**Head Flying Hours**		Time while head is positioning[*citation needed*]
240	0xF0	**Transfer Error Rate** (Fujitsu)		Count of times the link is reset during a data transfer.[24]
241	0xF1	**Total LBAs Written**		Total count of LBAs written
242	0xF2	**Total LBAs Read**		Total count of LBAs read Some S.M.A.R.T. utilities will report a negative number for the raw value since in reality it has 48 bits rather than 32.
250	0xFA	**Read Error Retry Rate**		Count of errors while reading from a disk
254	0xFE	**Free Fall Protection**		Count of "Free Fall Events" detected [25]

Threshold Exceeds Condition

Threshold Exceeds Condition (TEC) is an estimated date when a critical drive statistic attribute will reach its threshold value. When Drive Health software reports a "Nearest T.E.C.", it should be regarded as a "Failure date".

Prognosis of this date is based on the factor "Speed of attribute change"; how many points each month the value is decreasing/increasing. This factor is calculated automatically at any change of S.M.A.R.T. attributes for each attribute individually. Note that TEC dates are not guarantees; hard drives can and will either last much longer or fail much sooner than the date given by a TEC.

Self-tests

SMART drives may offer a number of self-tests:[26]

- `offline`
- `short`: usually under ten minutes
- `long`: extended: tens of minutes
- `conveyance`: `minutes`: identify damage incurred during transporting of the device

Selective self-tests of just a part of the surface may also be available.

See also

- Comparison of S.M.A.R.T. tools
- Predictive failure analysis

Notes

1. ^ **a** **b** Seagate statement on enhanced smart attributes (http://www.seagate.com/docs/pdf/whitepaper/enhanced_smart.pdf)
2. ^ Failure Trends in a Large Disk Drive Population (http://labs.google.com/papers/disk_failures.pdf) (Conclusion section), by Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz André Barroso, Google Inc. 1600 Amphitheatre Pkwy Mountain View, CA 94043
3. ^ **a** **b** PCTechGuide's page on S.M.A.R.T. (2003) (http://www.pctechguide.com/31HardDisk_SMART.htm)
4. ^ IBM Announcement Letter No. ZG92-0289 dated September 1, 1992 (http://www-306.ibm.com/common/ssi/OIX.wss?DocURL=http://d03xhttpcl001g.boulder.ibm.com/common/ssi/rep_ca/9/877/ENUSZG92-0289/index.html&InfoType=AN)
5. ^ http://web.archive.org/web/20080622210656/http://www.seagate.com/support/kb/disc/smart.html

6. ^ Compaq. IntelliSafe. Technical Report SSF-8035, Small Form Committee, January 1995.
7. ^ Stephens, Curtis E, ed. (December 11, 2006), "ATA/ATAPI Command Set (ATA8-ACS), working draft revision 3f" (http://www.t13.org/Documents/UploadedDocuments/docs2006/D1699r3f-ATA8-ACS.pdf) , *AT Attachment 8* (ANSI INCITS): pp. 198–213, 327–344, http://www.t13.org/Documents/UploadedDocuments/docs2006/D1699r3f-ATA8-ACS.pdf
8. ^ Hitachi Global Storage Technologies (19 September 2003), *Hard Disk Drive Specification: Hitachi Travelstar 80GN, revision 2.0* (http://www.hitachigst.com/tech/techlib.nsf/techdocs/85CC1FF9F3F11FE187256C4F0052E6B6/$file/80GNSpec2.0.pdf) , Hitachi Document Part Number S13K-1055-20, http://www.hitachigst.com/tech/techlib.nsf/techdocs/85CC1FF9F3F11FE187256C4F0052E6B6/$file/80GNSpec2.0.pdf
9. ^ Hatfield, Jim (September 30, 2005), *SMART Attribute Annex* (http://www.t13.org/Documents/UploadedDocuments/docs2005/e05148r0-ACS-SMARTAttributesAnnex.pdf) , e05148r0, http://www.t13.org/Documents/UploadedDocuments/docs2005/e05148r0-ACS-SMARTAttributesAnnex.pdf
10. ^ smartmontools FAQ ("Attribute 194 (Temperature Celsius) behaves strangely on my Seagate disk") (http://smartmontools.sourceforge.net/faq.html#temp-seagate)
11. ^ *Self-Monitoring, Analysis and Reporting Technology (SMART) :: Article* (http://smartlinux.sourceforge.net/smart/article.php) , 2009-03-10, http://smartlinux.sourceforge.net/smart/article.php
12. ^ Acronis Knowledge Base, 9109: S.M.A.R.T. Attribute: Power-On Hours (POH), http://kb.acronis.com/content/9109
13. ^ **a b c d e f** S.M.A.R.T. attribute list (ATA) (http://www.hdsentinel.com/help/en/56_attrib.html)
14. ^ *Fly Height Monitor Improves Hard Drive Reliability* (http://www.wdc.com/en/library/2579-850123.pdf) , Western Digital, April 1999, 79-850123-000, http://www.wdc.com/en/library/2579-850123.pdf
15. ^ **a b c** *Fujitsu MHT2080AT, MHT2060AT, MHT2040AT, MHT2030AT, MHG2020AT Disk Drives Product Manual* (http://www.fujitsu.com/downloads/COMP/fcpa/hdd/discontinued/mht20xxat_prod-manual.pdf) , Fujitsu Limited, 2003-07-04, C141-E192-02EN, http://www.fujitsu.com/downloads/COMP/fcpa/hdd/discontinued/mht20xxat_prod-manual.pdf
16. ^ laptop hard drive Load_Cycle_Count issue (http://ubuntuforums.org/showthread.php?p=5031046)
17. ^ Problem with hard drive clicking (http://www.thinkwiki.org/wiki/Problem_with_hard_drive_clicking) "Despite files being cached, POSIX-compliant file systems like ext2 or ext3 must update (=write) the last access time."
18. ^ Archlinux.org discussion list (http://bbs.archlinux.org/viewtopic.php?id=66706) "If linux tends to write to /var/log/* every 30s, then the heads can park/unpark every 30s."
19. ^ How to Reduce Power Consumption (http://www.thinkwiki.org/wiki/How_to_reduce_power_consumption#Hard_Drives) The files access time update, while mandated by POSIX, is causing lots of disks access; even accessing files on disk cache may wake the ATA or USB bus.
20. ^ Acronis Knowledge Base, 9132: S.M.A.R.T. Attribute: Reallocation Event Count, http://kb.acronis.com/content/9132
21. ^ Acronis Knowledge Base, 9133: S.M.A.R.T. Attribute: Current Pending Sector Count, http://kb.acronis.com/content/9133
22. ^ Lubomir Cabla (2009-08-06). "HDAT2 v4.6 User's Manual (Version 1.1)" (http://www.hdat2.com/files/hdat2en_v11.pdf) . http://www.hdat2.com/files/hdat2en_v11.pdf.
23. ^ "S.M.A.R.T. Linux project: Attributes" (http://smartlinux.sourceforge.net/smart/attributes.php) . http://smartlinux.sourceforge.net/smart/attributes.php.
24. ^ "MHY2xxxBH Disk Drives, Product/Maintenance Manual" (http://www.msc-ge.com/download/itmain/datasheets/fujitsu/MHY2xxxBH.pdf) . Fujitsu Limited. C141-E192-02EN. http://www.msc-ge.com/download/itmain/datasheets/fujitsu/MHY2xxxBH.pdf.
25. ^ Seagate Technology, LLC (September 2007), *Seagate Momentus 7200.2 SATA Product Manual, Publication Number: 100451238, Rev. D* (http://www.seagate.com/staticfiles/support/disc/manuals/notebook/momentus/7200.2/100451238d.pdf) , Hitachi Document Part Number S13K-1055-20, http://www.seagate.com/staticfiles/support/disc/manuals/notebook/momentus/7200.2/100451238d.pdf
26. ^ Manpage of SMARTCTL (http://smartmontools.sourceforge.net/man/smartctl.8.html) section on self-tests: "SMART RUN/ABORT OFFLINE TEST AND SELF-TEST OPTIONS: -t TEST, --test=TEST"

References

- "S.M.A.R.T. attribute meaning" (http://www.siguardian.com/products/siguardian/on_line_help/s_m_a_r_t_attribute_meaning.html) . *PalickSoft*. http://www.siguardian.com/products/siguardian/on_line_help/s_m_a_r_t_attribute_meaning.html. Retrieved February 3, 2006.
- Zbigniew Chlondowski. "S.M.A.R.T. Site: attributes reference table" (http://smartlinux.sourceforge.net/smart/attributes.php) . S.M.A.R.T. Linux. http://smartlinux.sourceforge.net/smart/attributes.php. Retrieved January 17, 2007.
- "S.M.A.R.T. attributes meaning" (http://www.ariolic.com/activesmart/smart-attributes/) . *Ariolic Software, Ltd*. 2007. http://www.ariolic.com/activesmart/smart-attributes/. Retrieved October 26, 2007.
- "Can we believe S.M.A.R.T. ?" (http://www.hdsentinel.com/smart) . *H.D.S. Hungary*. 2007. http://www.hdsentinel.com/smart. Retrieved June 4, 2008.
- Bruce Allen (2004). "Monitoring Hard Disks with SMART" (http://www.linuxjournal.com/article/6983) . Linux Journal. http://www.linuxjournal.com/article/6983. Retrieved August 8, 2010.

External links

- Out SMART Your Hard Drive (http://prefetch.net/articles/diskdrives.smart.html) Using the *smartmontools* program to monitor S.M.A.R.T. values
- How S.M.A.R.T. is your hard drive? (http://www.pc-king.co.uk/tips3.htm)
- How to predict hard disk failure (SMART Report) in Ubuntu Linux in just 3 clicks? (http://karuppuswamy.com/wordpress/2010/05/19/how-to-predict-hard-disk-failure-in-ubuntu-with-3-clicks/)