



SGI NAS

AutoSMART User Guide

v3.1.x

Copyright © 2013 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

SGI reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use.

Refer to the latest product announcement or contact your local SGI representative for information on feature and product availability.

This document includes the latest information available at the time of publication.

TRADEMARKS AND ATTRIBUTIONS

SGI, Silicon Graphics, Supportfolio and the SGI logo are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries. Solaris and OpenSolaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. Microsoft and Windows are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks mentioned herein are the property of their respective owners.

Document Number: 007-6306-001

Contents

Preface	v
1 Introduction to AutoSmart	1
About AutoSmart	1
Plugin Components	1
Self-Monitoring, Analysis, and Reporting Technology	2
2 Configuration	3
About AutoSMART Installation	3
Enabling AutoSMART	3
Disabling AutoSMART	4
Checking the Status of the Disks	5
Managing SMART Collector	5
Enabling the Autosmart Fault Trigger	6
Viewing Graphs	6
A ATA/SATA Attribute Descriptions	9
B SCSI SMART Attributes	15
C SMART States	17
Index	19

Preface

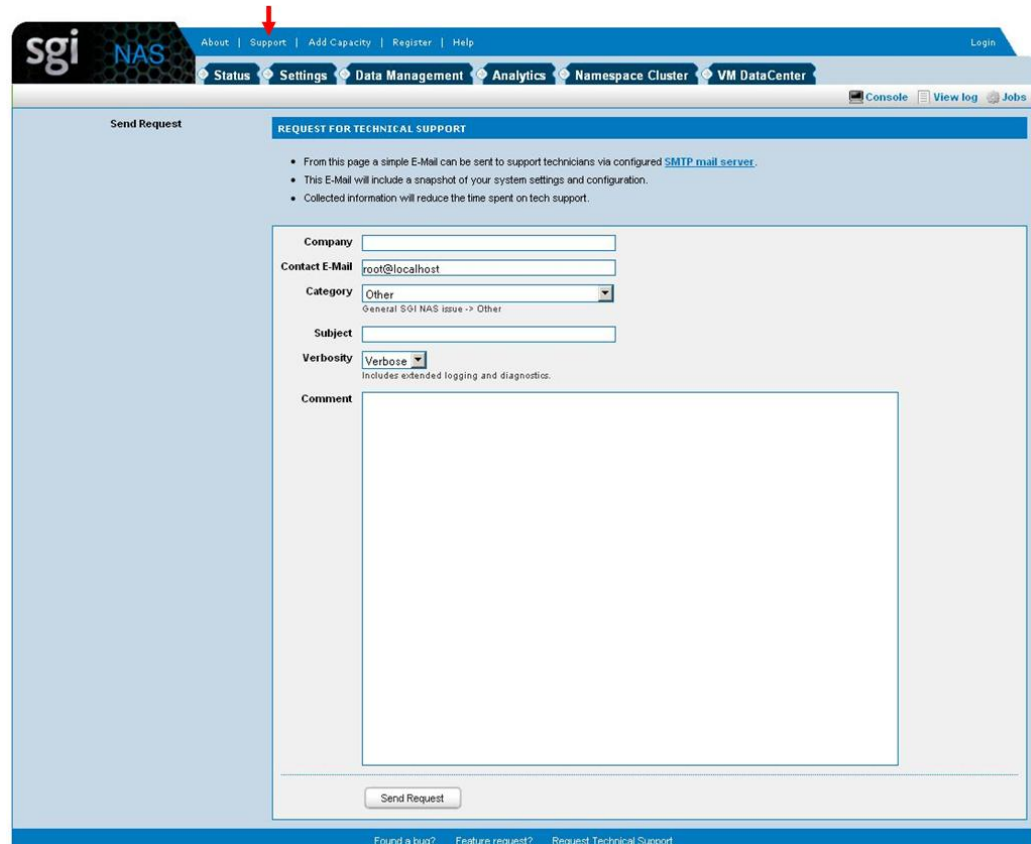
This documentation presents information specific to SGI NAS products. The information is for reference purposes and is subject to change.

Intended Audience

This documentation is intended for Network Storage Administrators and assumes that you have experience with data storage concepts, such as NAS, SAN, NFS, and ZFS.

Contacting Support

To contact support at SGI, click the Support link in NMV as marked with a red arrow on the screen below:



The screenshot shows the SGI NAS web interface. At the top, there is a navigation bar with the SGI logo and 'NAS' text. The navigation bar includes links for 'About', 'Support', 'Add Capacity', 'Register', and 'Help'. A red arrow points to the 'Support' link. Below the navigation bar, there are tabs for 'Status', 'Settings', 'Data Management', 'Analytics', 'Namespace Cluster', and 'VM DataCenter'. The main content area is titled 'Send Request' and contains a 'REQUEST FOR TECHNICAL SUPPORT' form. The form includes the following fields and options:

- Company:** Text input field.
- Contact E-Mail:** Text input field with the value 'root@localhost'.
- Category:** Dropdown menu with 'Other' selected. Below it, the text 'General SGI NAS issue -> Other' is visible.
- Subject:** Text input field.
- Verbosity:** Dropdown menu with 'Verbose' selected. Below it, the text 'Includes extended logging and diagnostics' is visible.
- Comment:** Large text area for providing details.

At the bottom of the form is a 'Send Request' button. Below the form, there are links for 'Found a bug?', 'Feature request?', and 'Request Technical Support'.

or type the following NMC command:

```
nmc:/$ support
```

which will then prompt for a subject and message.

Other resources

For licensing questions, please contact your SGI sales or support representative.

Product Support

SGI provides a comprehensive product support and maintenance program for its products.

For a full description of this program, do one of the following:

- See <http://www.sgi.com/support/>.
- If you are in North America, contact the Technical Assistance Center at 1 (800) 800 4SGI or contact your authorized service provider.
- If you are outside North America, see the following website for the appropriate Customer Service phone number: <http://www.sgi.com/support/supportcenters.html>.

Introduction to AutoSmart

This section includes the following topics:

- [About AutoSmart](#)
- [Plugin Components](#)
- [Self-Monitoring, Analysis, and Reporting Technology](#)

About AutoSmart

The AutoSMART (Self Monitoring Analysis and Reporting Technology) plugin uses technology that provides analysis and diagnostics of hard disk drives. Therefore, it decreases the risk of information loss on your SGI NAS appliance. It provides statistical information in textual, numerical and graph variants.

ATA (or SATA) drives and SCSI (SAS) drives typically use different messaging systems. ATA drives use **Self-Monitoring, Analysis, and Reporting Technology** (SMART), while SCSI drives use Informational Exceptions (IE) to report faults. The AutoSMART plugin collects information from both types of disks. The term SMART is used in the context of disk hardware monitoring without considering the type of disk.

Plugin Components

The plugin consists of the following components:

- Smart-collector
- Smart Settings page
- Graph page
- NMC plugin

Smart-collector periodically gathers SMART statistics and status information from hard drives. By default it is set to one hour. If any of the monitored hard drives fail, it sends a fault trigger notification within the system and by e-mail.

Self-Monitoring, Analysis, and Reporting Technology

S.M.A.R.T. is the technology of hard disk drive status checking and fault prognosis by monitoring certain online hard drive activities. It was developed in order to provide higher reliability for data storage.

SMART technology works by monitoring the hard drive's SMART parameters independently. If AutoSMART detects an extreme value, it sends a warning message reporting device pre-emergency status. When the User receives the message, they have an opportunity to perform the necessary preventive operations.

AutoSMART divides failures into two categories:

- **Predictable**

Predictable failures are characterized by step-by-step degradation of one or another parameter. Many mechanical failures are usually regarded as predictable. With the development of SMART technology, the increasing number of failures becomes predictable, thereby increasing the probability of the failure prediction. While disks drives run, SMART monitors all of the occurring errors and the suspicious phenomena which are reflected in appropriate attributes.

- **Unpredictable**

Unpredictable failures occur fast, such as from a voltage surge damaging the electronic circuit, from damage to the magnetic heads and from impact damage to the magnetic disk surface. Improvement of quality, construction, technology and production reduces the probability of unpredictable crashes.

The quantity of attributes for various models and manufacturers differs. Each attribute has a value from 0 to 253. The value is the reliability of the concrete attributes and standard values defined by the manufacturer. A high value means the absence of changes of the given parameter. Low value means fast degradation or a possible forthcoming failure. Therefore, the higher the value of the attribute, the better.

For each attribute, the manufacturer guarantees the minimum possible value for operation of the disk drive. When the value of the attribute is lower than the guaranteed success value, the failure in operation or total failure is very probable.

Attributes can be critical and noncritical:

- **Critical**

Parameter performs outside of the normal limits. This indicates a failure.

- **Noncritical**

Parameter performs outside of the normal limits. This indicates that there is a problem at which the hard disk drive preserves its working capacity, though with some performance degradation.

Configuration

This section includes the following topics:

- [About AutoSMART Installation](#)
- [Enabling AutoSMART](#)
- [Disabling AutoSMART](#)
- [Checking the Status of the Disks](#)
- [Managing SMART Collector](#)
- [Enabling the Autosmart Fault Trigger](#)
- [Viewing Graphs](#)

About AutoSMART Installation

The AutoSMART plugin installs automatically when you install SGI NAS.

- ❖ *To view plugins, using NMC:*
 - ◆ Type:

```
nmc:/$ show plugin
```
- ❖ *To view plugins, using NMV:*
 1. Click **Settings > Appliance**.
 2. In the Administration panel, click **Plugins**.

Enabling AutoSMART

You can enable or disable the smart statistic data collection for all disks or specific disks.

- ❖ *To enable AutoSMART for all or specific disks, using NMV:*
 1. Click **Settings > Disks**.
 2. In the Smart panel, click **Smart Settings**.

- 3. In the SMART enable column, select the disks for which to enable the SMART settings.
- ❖ *To enable AutoSMART for specific disks, using NMC:*
 - ◆ Type:


```
nmc:/$ setup lun smart enable -d <disk_name>
```

 Example:


```
nmc:/$ setup lun smart enable -d c1t3d0
```
- ❖ *To enable AutoSMART for all disks, using NMC:*
 - ◆ Type:


```
nmc:/$ setup lun smart enable
```

 System response:


```
Enable SMART monitoring          : (Use SPACEBAR for multiple
selection)
c1t0d0  c1t1d0  c1t2d0  c1t3d0
-----
Enable SMART monitoring for selected disks. Navigate with
arrow keys (or hjkl), or Ctrl-C to exit.
```

Note: There is a limitation on how many devices can be displayed.

Disabling AutoSMART

You can exclude all disks from SMART monitoring or you can exclude specific disks.

- ❖ *To enable AutoSMART for all or specific disks, using NMV:*
 - 1. Click **Settings > Disks**.
 - 2. In the SMART panel, click **Smart Settings**.
 - 3. In the SMART enable column, deselect the disks for which to disable the SMART settings.
- ❖ *To disable AutoSMART for a specific disk, using NMC:*
 - ◆ Type:


```
nmc:/$ setup lun smart disable -d <disk_name>
```

 Example:


```
nmc:/$ setup lun smart disable -d c1t3d0
```

- ❖ *To disable AutoSMART for all disks:*

- ◆ Type:

```
nmc:/$ setup lun smart disable
```

System response:

```
Disable SMART monitoring      : (Use SPACEBAR for multiple
selection)
```

```
c1t0d0  c1t1d0  c1t2d0  c1t3d0
```

```
-----
Disable SMART monitoring for selected disks. Navigate with
arrow keys (or hjkl), or Ctrl-C to exit.
```

Note: There is a limitation on how many devices can be displayed.

Checking the Status of the Disks

You can check the status of the disks at any time.

- ❖ *To check the status of the AutoSMART plugin, using NMC:*

- ◆ Type:

```
nmc:/$ show lun smartstat
```



Note:

During the process of initial SMART data collecting, the output of **show lun smartstat** may display that a disk is in a Failed state. This does not indicate actual disk failure.

For more information about SMART states of the disk, see Appendix C. SMART states.

Managing SMART Collector

Manage the Smart Collector using the smart-collector page.

- ❖ *To enable SMART Collector, using NMV:*

1. Click **Data Management > Runners**.
2. In the Registered Statistics Collectors panel, click **smart-collector**.
3. Set Status to **Enabled**.
4. Set the Trace level.
 - **Disabled**
Disables SMART collector.

- **Logging**
Displays the log files.
 - **Verbose**
Displays the corresponding values, in full detail.
 - **Debug**
Displays debug information for development.
5. Set the Periodic Interval. (Default is set to one hour.)
 6. Click **Save** or **Run Now**.



Note:

SMART switches off automatically for virtual disks, because they are not supported. Virtual disk SMART state always displays as **N/A-UNSUPP**.

Enabling the Autosmart Fault Trigger

The Autosmartcheck Fault Trigger checks the state of the SMART drives and reports if the SMART subsystem detects that certain hard disks are generating errors.

❖ *To enable the Autosmartcheck Fault Trigger, using NMV:*

1. Click **Data Management > Runners**.
2. In the Registered Fault Triggers panel, click **nms-autosmartcheck**.
3. Set Status to **Enabled**.
4. Set the Trace level.
 - **Disabled**
Disables SMART collector.
 - **Logging**
Displays the log files.
 - **Verbose**
Displays the corresponding values, in full detail.
 - **Debug**
Displays debug information for development.
5. Set the Periodic Interval. (Default is set to one hour.)
6. Click **Save** or **Run Now**.

Viewing Graphs

You can view the autosmart data on a graph.

- ❖ *To view the autosmart data in graphical format, using NMV:*
 1. Click **Status > Storage**.
 2. In the Volume Statistics panel, click **AutoSMART statistics**.
 3. Select a relevant disk for which to view the statistics.
 4. Select the relevant parameter.
 5. Select the beginning and end times for the data.
 6. Click **Redraw**.



ATA/SATA Attribute Descriptions

The following table describe the attributes of ATA/SATA disk drives.



Access the following data by clicking on **Settings > Disks** and clicking on the disk after enabling SMART settings.

Table A-1: ATA/SATA Attributes Description

ID	Attribute Name	Description
01	Read Error Rate	Indicates the rate of hardware read errors that occurred when reading data from a disk surface. The raw value has a different structure for different vendors and is often not meaningful as a decimal number.
02	Throughput Performance	Overall (general) throughput performance of a hard disk drive. If the value of this attribute decreases, it indicates that there is a problem with the disk.
03	Spin-Up Time	Average time of spindle spin up (from zero RPM to fully operational [milliseconds]).
04	Start/Stop Count	Tallies spindle start/stop cycles. The spindle turns on, the count increases, both when the hard disk turns on after having been turned entirely off, (disconnected from power source) and when the hard disk returns from sleep mode.
05	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 called remapping. Reallocated sectors are called remaps. This is why, on modern hard disks, you cannot find bad blocks while testing the surface; all bad blocks are hidden in reallocated sectors. However, as the number of reallocated sectors increases, the read/write speed tends to decrease. The raw value normally represents a count of the number 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	Read Channel Margin	Margin of a channel while reading data. The function of this attribute is not specified.

Table A-1: ATA/SATA Attributes Description (Continued)

ID	Attribute Name	Description
07	Seek Error Rate	Rate of seek errors of the magnetic heads. If there is a partial failure in the mechanical positioning system, then seek errors occur. 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 is structured differently for different vendors and is often not meaningful as a decimal number.
08	Seek Time Performance	Average performance of seek operations of the magnetic heads. If this attribute decreases, it is a sign of problems in the mechanical subsystem.
09	Power-On Hours (POH)	Count of hours in the power-on state. The raw value of this attribute displays total count of hours (or minutes, or seconds, depending on the manufacturer) in the power-on state.
10	Spin Retry Count	Count of retries of the spin start attempts. This attribute stores a total count of the spin start attempts to reach the fully operational speed (assuming that the first attempt was unsuccessful). An increase of this attribute value is a sign of problems in the hard disk mechanical subsystem.
11	Recalibration Retries Calibration Retry Count	Indicates the number of times recalibration was requested (assuming that the first attempt was unsuccessful). An increase of this attribute value is a sign of problems in the hard disk mechanical subsystem.
12	Power Cycle Count	Indicates the count of full hard disk power on/off cycles.
13	Soft Read Error Rate	Uncorrected read errors reported to the operating system.
183	SATA Downshift Error Count	Western Digital and Samsung attribute.
185	Head Stability	Western Digital attribute.
186	Induced Op-Vibration Detection	Western Digital attribute.
187	Reported Uncorrectable Errors	The number of errors that could not be recovered using hardware ECC (see attribute 195).
188	Command Timeout	The number of aborted operations due to HDD timeout. Normally this attribute value is equal to zero. If the value is higher than zero, then most likely there will be serious problems with the power supply or an oxidized data cable.

Table A-1: ATA/SATA Attributes Description (Continued)

ID	Attribute Name	Description
189	High Fly Writes	<p>Indicates the count of these errors detected over the lifetime of the drive.</p> <p>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 it encounters an unsafe fly height condition, the write process stops, and the information is rewritten or reallocated to a safe region of the hard drive.</p> <p>This feature is implemented in most modern Seagate drives and some of Western Digital's drives, starting with the WD Enterprise WDE18300 and the WDE9180 Ultra2 SCSI hard drive.</p>
190	Airflow Temperature (WDC)	Airflow temperature on Western Digital HDs.
190	Temperature Difference from 100	Value is equal to $(100 - \text{temp. } ^\circ\text{C})$, allowing a manufacturer to set a minimum threshold which corresponds to a maximum temperature.
191	G-sense Error Rate	Number of errors resulting from externally-induced shock & vibration.
192	Power-off Retract	Number of times the heads are loaded off the media.
	Count Emergency Retract Cycle count	Number of times the heads can be unloaded without actually powering off.
193	Load Cycle Count Load/Unload Cycle Count (Fujitsu)	<p>Count of load/unload cycles into head landing zone position.</p> <p>The typical lifetime rating for laptop (2.5-in) hard drives is 300,000 to 600,000 load cycles. Some laptop drives are programmed to unload the heads after five seconds of inactivity. Many Linux installations write to the filesystem a few times a minute in the background. As a result, there may be 100 or more load cycles per hour, and you may exceed the load cycle rating in less than a year.</p>
194	Temperature	Current internal temperature.
195	Hardware ECC Recovered	The raw value has different structure for different vendors and is often not meaningful as a decimal number.
196	Reallocation Event Count	Count of remap operations. The raw value of this attribute displays the total number of attempts to transfer data from reallocated sectors to a spare area. It counts both successful & unsuccessful attempts.
197	Current Pending Sector Count	Number of unstable (waiting to be remapped because of read errors) sectors. If an unstable sector is subsequently written or read successfully, this value decreases and the sector is not remapped. Read errors on a sector do not remap the sector (since it might be readable later); instead, the drive firmware remembers that the sector needs remapping, and remaps it the next time it is written.

Table A-1: ATA/SATA Attributes Description (Continued)

ID	Attribute Name	Description
198	Uncorrectable Sector Count	The total number 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, (Off-Line Scan Uncorrectable Sector Count – Fujitsu).
199	UltraDMA CRC Error Count	Number of errors in data transfer through the interface cable as determined by ICRC (Interface Cyclic Redundancy Check).
200	Multi-Zone Error Rate	Total number of errors when writing to multiple sectors.
200	Write Error Rate	Total number of errors when writing to a sector.
201	Soft Read Error Rate	Number of off-track errors.
202	Data Address Mark errors	Number of Data Address Mark errors (or vendor-specific).
203	Run Out Cancel	Number of ECC errors.
204	Soft ECC Correction	Number of errors corrected by software ECC.
205	Thermal Asperity Rate (TAR)	Number of errors due to high temperature.
206	Flying Height	Height of heads above the disk surface. A flying height that is too low increases the chances of a head crash while a flying height that is too high increases the chances of a read/write error.
207	Spin High Current	Amount of surge current used to spin up the drive.
208	Spin Buzz	Number of buzz routines needed to spin up the drive due to insufficient power.
209	Offline Seek Performance	Drive's seek performance during its internal tests.
211	Vibration During Write	Indicates vibrations encountered during write operations.
212	Shock During Write	Indicates shocks encountered during write operations.
220	Disk Shift	Distance the disk has shifted relative to the spindle, (usually due to shock or temperature). Unit of measure is unknown.
221	G-Sense Error Rate	Number of errors resulting from externally-induced shock & vibration.
222	Loaded Hours	Time spent operating under data load, (movement of magnetic head armature).
223	Load/Unload Retry Count	Number of times head changes position.
224	Load Friction	Resistance caused by friction in mechanical parts while operating.
225	Load/Unload Cycle Count	Total number of load cycles
226	Load In-time	Total time of loading on the magnetic heads actuator (time not spent in parking area).
227	Torque Amplification Count	Number of attempts to compensate for platter speed variations.

Table A-1: ATA/SATA Attributes Description (Continued)

ID	Attribute Name	Description
228	Power-Off Retract Cycle	Number of times the magnetic armature was retracted automatically as a result of cutting power.
230	GMR Head Amplitude	Amplitude of thrashing (distance of repetitive forward/reverse head motion).
231	Temperature	Drive temperature.
240	Head Flying Hours	Duration for head positioning.
240	Transfer Error Rate	Counts the number of times the link resets during data transfer.
241	Total LBAs Written	Total logical block addressing units written. Although degradation of this parameter can be an indicator of drive aging and/or potential electromechanical problems, it does not directly indicate imminent drive failure.
242	Total LBAs Read	Some SMART utilities report a negative number for the raw value since it has 48 bits rather than 32.
250	Read Error Retry Rate	Number of errors while reading from a disk.
254	Free Fall Protection	Number of Free Fall events detected.

SCSI SMART Attributes

The following table describes SCSI SMART attributes and contains the following classifications:

- **Absolute**
Measures absolute value.
- **Counter**
Counts number of events from the start of disk use.
- **Per hour**
Counts number of events per hour. Smart-collector runner can run for any period of time. It automatically recalculates the value to an hourly rate.



Note:

Access the following data by clicking on **Settings > Disks** and clicking on the disk after enabling SMART settings.

Table B-1: SCSI SMART Attributes

Attribute	Description	Type
Blocks read from cache and sent to initiator		per hour
Blocks received from initiator		per hour
Blocks sent to initiator		per hour
Current Drive Temperature	Temperature of drive in degrees Celsius.	absolute
Delayed ECC corrections during reading	An error code or algorithm is applied in order to obtain perfect data with a substantial delay. This means that the correction took longer than a sector time so that reading/writing of subsequent sectors was delayed.	per hour
Delayed ECC corrections during writing	See Delayed ECC corrections during reading.	per hour
Elements in grown defect list	Displays how many blocks defined as 'bad blocks' for a defined period of time.	per hour

Table B-1: SCSI SMART Attributes

Attribute	Description	Type
Fast ECC corrections during reading	An error correction was applied to obtain perfect data for a period of time. The correction did not postpone reading the later sectors (e.g. a revolution was not lost). The counter is incremented once for each logical block that requires a correction.	per hour
Fast ECC corrections during writing	See Fast ECC corrections during reading.	per hour
Gigabytes read processed	Specifies the counter that counts the number of bytes between two data collections, either successfully or unsuccessfully read, written or verified (depending on the log page) from the drive. If a transfer terminates early because of an unrecoverable error, it counts only the logical blocks up to and including the one with the uncorrected data.	per hour
Gigabytes write processed	See Gigabytes read processed.	per hour
Non-medium error count	Represents the number of recoverable events other than write, read or verify errors.	per hour
Number of hours powered up		per hour
Rereads	Specifies the counter during periods of time, and counting the number of errors that are corrected by applying retries. This counts errors recovered, not the number of retries. If five retries were required to recover one block of data, the counter increments by one, not five. The counter is incremented once for each logical block that is recovered using retries	per hour
Rewrites	See Rereads.	per hour
Total Uncorrected read errors	Specifies the counter that contains the total number of blocks for which an uncorrected data error has occurred.	counter
Total Uncorrected write errors	Specifies the counter that contains the total number of blocks for which an uncorrected data error has occurred.	counter



SMART States

The following table describe the SMART states.



Access the following data by clicking on **Settings > Disks** and clicking on the disk after enabling SMART settings.

Table C-1: SMART States

State	Short Description	Long Descriptions
FAILED	Disk faulted.	Disk failure, or out of order imminent.
N/A	Not available	State not defined.
N/A-LOGICAL	Not available for logical disks.	Autosmart info not available for virtual and logical disks.
N/A-SPARED	Not available for disks spared by zfs.	Autosmart info not available for disks defined, as faulted, by a zfs diagnostic system.
N/A-SSD	Not available for SSD	SSD disks SMART info unavailable.
N/A-UNSUPP	Not available for unsupported controller-disk.	Cannot obtain any SMART info for current disk.
N/A-WARNING	Not available for warned disks.	Attempts to obtain SMART info issues warnings and faults
U/M	Unmanaged	Autosmart disabled for current disk by user.

Index

A

- ATA/SATA
 - attribute descriptions 9
- attributes
 - ATA/SATA 9
 - critical 2
 - noncritical 2
 - SCSI drives 15
- autoSMART
 - about 1
 - disabling 4
 - enabling 3
 - installation 3

C

- checking status
 - disks 5
- components
 - plugin 1
- critical attributes 2

D

- disabling
 - autoSMART 4
- disks
 - checking status 5

E

- enabling
 - autoSMART 3
 - fault trigger 6

F

- failures
 - predictable 2
 - unpredictable 2
- fault trigger
 - enabling autoSMART 6

G

- graphs
 - viewing 6
- guide
 - audience v

I

- installation

- autoSMART 3

M

- managing SMART
 - collector 5

N

- noncritical attributes 2

P

- plugin
 - components 1
- predictable failures 2

S

- self-monitoring, analysis,
 - and reporting
 - technology 2
- SMART collector
 - managing 5
- SMART states 17
- support
 - contact v

U

- unpredictable failures 2

V

- viewing
 - graphs 6