



Questions and Answers about HDD

Version 1.0
October 2011

Copyright

Copyright@2004~2011, Qsan Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from Qsan Technology, Inc.

Trademarks

All products and trade names used in this manual are trademarks or registered trademarks of their respective companies.

Qsan Technology, Inc.

2F., No.23, Lane 583, Ruiguang Rd.
Neihu Dist., Taipei 114
Taiwan, R.O.C.

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

Email: sales@Qsan.com.tw
Website: www.QsanTechnology.com

Seagate Hard Drive

Question:

When we use new Seagate nearline SAS 2TB hard drives, we are experiencing slow initialization on creating RAID 5 or RAID 6.

Answer:

This cases happens on brand-new Seagate Enterprise hard drives. Seagate answers that Enterprise HDD will enable the feature of Read After Write (RAW) when the POH (Power On Hour) is less than 24 hours to make sure the data integrity. That's why the slow initialization happens on these new drives in the beginning, but the speed will become normal after one or two days.

We purchased new Seagate Constellation ES, ST32000444SS hard drives and tested by lometer on the server directly. The write performace is 4MB/s only, but normally is 130MB/s after using for a while.

This feature can not be disable. The workaround solution is to power on the system with these new drives and the performance should become normal after one or two days.

Question:

When using Seagate Constellation ES ST2000NM0011 SATA 6Gb/s on P160Vxx / P160C, why the standby of hard drives do not work?

Answer:

Seagate answers that Seagate Constellation ES hard drive will enable standby longer than 3 minutes. So it has to set standby 5 minutes or 30 minutes on P160V Series / P160C.

3TB / 3G MUX / 6G MUX

Question:

How to support 3TB HDD and SATA 6G/s?

Answer:

Arrange the tables on the following. Be aware the FW version and the comments.

Dual controllers with LSI 3G expander (System / Controller):

- F300Q-Dxxx / F300H (FW 2.1.1, Local Expander FW 2.1.0)
- F400Q-Dxxx / F400H (FW 2.1.1, Local Expander FW 2.1.0)
- P300Q-Dxxx / P300H (FW 2.1.1, Local Expander FW 2.1.0)
- P500Q-Dxxx / P500H (FW 2.1.1, Local Expander FW 2.1.0)
- S300Q-Dxxx / S300H (FW 2.1.1, Local Expander FW 2.1.0)

HDD Type		6G MUX board	3G MUX board
SATA	3TB	2794 GB	2048 GB (*)
	SATA 6Gb/s	SATA 3Gb/s	SATA 3Gb/s
	SATA 3Gb/s	SATA 3Gb/s	SATA 3Gb/s
	SATA 1.5Gb/s	SATA 3Gb/s	N/A (**)
Without MUX board			
SAS	3TB	2794 GB	
	SAS 6Gb/s	SAS 3Gb/s	
	SAS 3Gb/s	SAS 3Gb/s	

- (*) Dual controllers with LSI 3G expander has 2TB limitation for SATA 3TB hard drive. It has to connect with 6G MUX board to broke 2TB.
- (**) Dual controllers with LSI 3G expander does not support SATA 1.5Gb/s with 3G MUX board.

Dual controllers with LSI 6G expander (System / Controller):

- F600Q-Dxxx / F600H (FW 1.0.0)
- P400Q-Dxxx / P400H (FW 1.0.0)
- P600Q-Dxxx / P600H (FW 1.0.0)

HDD Type		6G MUX board	3G MUX board
SATA	3TB	2794 GB	2794 GB
	SATA 6Gb/s	SATA 6Gb/s	SATA 3Gb/s
	SATA 3Gb/s	SATA 6Gb/s	SATA 3Gb/s
	SATA 1.5Gb/s	SATA 6Gb/s	SATA 1.5Gb/s
Without MUX board			
SAS	3TB	2794 GB	
	SAS 6Gb/s	SAS 6Gb/s	
	SAS 3Gb/s	SAS 3Gb/s	

Single controller with LSI 3G expander (System / Controller):

- F300Q-Cxxx / F300C (FW 2.1.1, Local Expander FW 2.1.0)
- P300Q-Cxxx / P300C (FW 2.1.1, Local Expander FW 2.1.0)
- P500Q-Cxxx / P500C (FW 2.1.1, Local Expander FW 2.1.0)
- S300Q-Cxxx / S300C (FW 2.1.1, Local Expander FW 2.1.0)

- F200Vxx / F200C FW (FW 1.0.11, Local Expander FW 2.1.0)
- P210Vxx / P210C FW (FW 1.0.11, Local Expander FW 2.1.0)
- S510Vxx / S510C FW (FW 1.0.11, Local Expander FW 2.1.0)

HDD Type		6G MUX board	Without MUX board
SATA	3TB	2794 GB	2048 GB (*)
	SATA 6Gb/s	SATA 3Gb/s	SATA 3Gb/s
	SATA 3Gb/s	SATA 3Gb/s	SATA 3Gb/s
	SATA 1.5Gb/s	SATA 3Gb/s	N/A (**)
			Without MUX board
SAS	3TB	2794 GB	
	SAS 6Gb/s	SAS 6Gb/s	
	SAS 3Gb/s	SAS 3Gb/s	

- (*) Single controller with LSI 3G expander has 2TB limitation for SATA 3TB hard drive. It has to connect with 6G MUX board to broke 2TB.
- (**) Single controller with LSI 3G expander does not support SATA 1.5Gb/s without MUX board

Single controller with Marvell SATA controller (System / Controller):

- P110Vxx / P110C / P130C (FW 1.0.9)
- P160Vxx / P160C (FW 1.0.9)
- S121Vxx / S121C (FW 1.0.9)

HDD Type		Without MUX board
SATA	3TB	2794 GB
	SATA 6Gb/s	SATA 3Gb/s
	SATA 3Gb/s	SATA 3Gb/s
	SATA 1.5Gb/s	SATA 1.5Gb/s
		Without MUX board
SAS	3TB	N/A (*)
	SAS 6Gb/s	N/A (*)
	SAS 3Gb/s	N/A (*)

- (*) Single controller with Marvell SATA controller does not support SAS HDD.

S.M.A.R.T.

Question:

What is S.M.A.R.T. and what is its usage?

Answer:

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

Moreover, this technology is intended to recognize the conditions that indicate imminent drive failure and is designed to provide sufficient warning of a failure to allow user to back up the data before an actual failure occurs.

Question:

How does S.M.A.R.T. work?

Answer:

Using the built-in commands in S.M.A.R.T. aware hard drive and the monitoring tools in host can analyze and compare the magnetic head, motor, disks, circuit operating status, history records, and default thresholds of hard drive. Once the threshold is reached, it will alert users automatically.

Question:

Do SAS hard drives have S.M.A.R.T. attributes?

Answer:

Only ATA/SATA/SCSI hard drives have S.M.A.R.T. values. SAS hard drives have their own mechanisms for monitoring.

Question:

When the event log displays S.M.A.R.T. error in Qsan system, does it mean that the physical drive is ill-conditioned?

Answer:

S.M.A.R.T. values display the tendency of hard drive failure. An ill-conditioned hard drive does not mean that its S.M.A.R.T. would have error. In addition, Qsan system reads the disk S.M.A.R.T. values which are provided by disks themselves, and displays them on web UI. They are only for

reference and the meaning of the values are provided by disk vendors. QSAN system does not judge the disk health by S.M.A.R.T. attributes. Qsan system diagnoses the disk health by real I/O, such as I/O read/write error, I/O time out ...etc.

Question:

What are the meanings of the following S.M.A.R.T. attributes on web UI?

1. Read error rate
2. Spin up time
3. Reallocated sector count
4. Seek error rate
5. Spin up retries
6. Calibration retries
7. Temperature

Answer:

Qsan selects some important S.M.A.R.T. attributes and put them on web UI. Please refer the reference of Wikipedia - S.M.A.R.T..

Attribute ID	Attribute name
01	Read Error Rate
03	Spin-up Time
05	Reallocated Sector Count
07	Seek Error Rate
10	Spin Up Retry count
11	(Re)calibration Retries
194	Temperature

The detailed descriptions are listed as follows:

1. Read error rate (lower raw value is better)

Indicates the rate of hardware read errors that occurred when reading data from a disk surface. A non-zero value indicates a problem with either the disk surface or read/write heads. Note that Seagate drives often report a raw value that is very high even on new drives, and does not thereby indicate a failure.
2. Spin up time (lower raw value is better)

Average time of spindle spin up (from zero RPM to fully operational [millisecs]).

3. Reallocate sector count (lower raw value is better)
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. This is why, on modern hard disks, "bad blocks" cannot be found 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.
4. Seek error rate (lower raw value is better)
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. More seek errors indicates a worsening condition of a disk's surface or the mechanical subsystem, or both. Note that Seagate drives often report a raw value that is very high, even on new drives, and this does not normally indicate a failure.
5. Spin up retries (lower raw value is better)
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.
6. Calibration retries (lower raw value is better)
This attribute indicates the number of times 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.
7. Temperature (lower raw value is better)
Current internal temperature.

Take an example of "Read error rate: 200(51)", the value 200 is normalized value and the value 51 in parentheses is threshold value. All these values are defined by the disk vendors and varied among different vendors. Even different models made by same vendor may have different definition of S.M.A.R.T.. The normalized value is just like a score and higher score is better. But the threshold value is the bottom line of the attribute. When the normalized value is under the threshold value, it means that this disk might fail at any time.

Question:

For example, if the event log displays "S.M.A.R.T. Threshold Exceed Condition occurred for attribute reallocated sector count", what user should do with that hard drive?

Answer:

It means that the hard drive is going to be broken. The data might be lost if the use continue using this hard drive. Because bad sector remap pool of that hard drive is going to be run out. Please change another well-conditioned hard drive instead.

Question:

If S.M.A.R.T. reports some serious problems with a hard disk, should I replace it? (Cite from <http://smartlinux.sourceforge.net/smart/faq.php?#11>)

Answer:

S.M.A.R.T. is designed to detect upcoming failures in the hard disk drive. While it is possible that the hard disk drive may not fail for several months it is still possible that the hard disk drive may fail shortly. It is ok to continue to use the hard disk drive, however we highly recommend that if you are still under warranty, you should consider getting the drive replaced before your warranty runs out.

Reference

- Hard Drives | Seagate
<http://www.seagate.com>
- Wikipedia - S.M.A.R.T.
<http://en.wikipedia.org/wiki/S.M.A.R.T.>
- S.M.A.R.T. :: FAQ
<http://smartlinux.sourceforge.net/smart/faq.php>
- How Hard Disk S.M.A.R.T. Works?
<http://www.hdsentinel.com/smart/index.php>
- S.M.A.R.T. attribute list
<http://www.hdsentinel.com/smart/smartattr.php>