

SCSI TOOLBOX, LLC

Disk Drive Testing 101

Contents

Introduction 3

Is the drive working or not? 3

Can the drive tell us it has had problems?..... 5

 SCSI/SAS/FC Log Pages..... 5

 ATA/SATA SMART Information 7

Can the drive reliably write and read data? 8

Is the drive set up appropriately? 9

What is the performance profile of your drive? 10

Summary 13

Introduction

Disk drives are complex devices, marvels of mechanical engineering and real-time computing magic.

For example, the heads of a rotating magnetic disk physically fly over the moving platters, flying at a height of as little as 3 nanometers at a speed over 128 mph! No wonder bumps and drops can do so much damage – just like flying an extremely fragile airplane into metal-hard ground at 128 mph! Not only can the airplane (drive head) be damaged, but the ground (platter surface) can be dug-up, furloughed, damaged. The magnetic coating on the platter is like a thin layer of top soil – scraping it away scrapes away data.

Bottom line – be gentle with your disk drives. You may be under pressure to get a large number of drives tested by the end of the day – but take your time. Move them gently and slowly. **Never** move a drive while it is spinning. Never drop a drive or bump it into anything. And keep the drive cooled while testing – **never** power up a disk drive without some kind of fan to move air around the drive to draw away the heat it generates.

All disk drives have a built-in computer to control all the physical operations of the drive as well as dealing with data encoding/decoding, queuing, transferring data in and out – a marvel of real-time computing!

Disk drives can work perfectly, or they cannot work at all, or they can “sort of” work, work marginally or poorly. The goal of basic disk drive testing is to:

1. determine if the drive is working at all or not
2. determine if the drive itself can tell us if it has had a problem in the past
3. determine if the drive can reliably store and retrieve data
4. determine if drive settings are appropriate for the intended use of the drive and,
5. determine the performance characteristics of the drive.

Is the drive working or not?

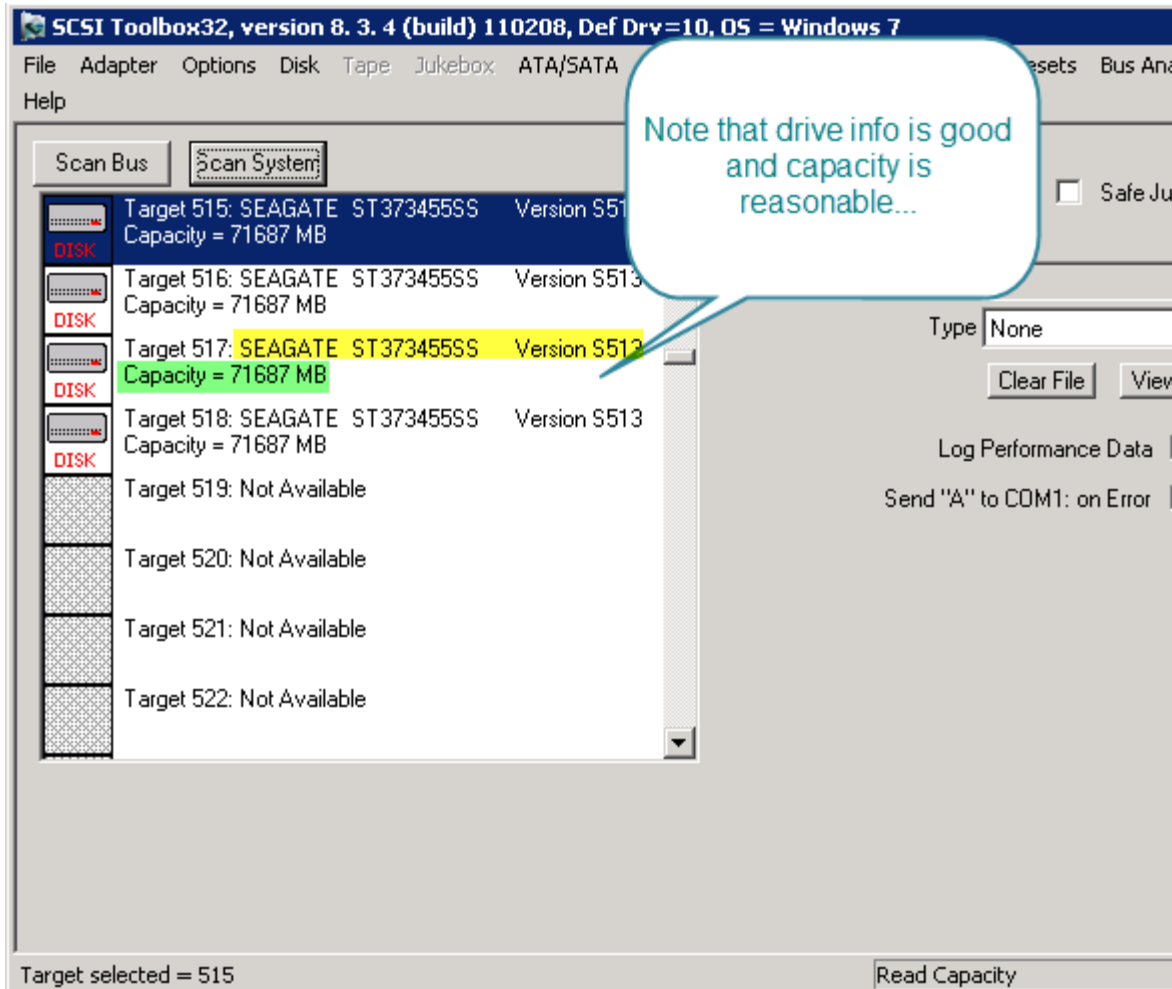
Determining if the drive is alive or not is simply a matter of connecting it to a test system and checking that the drive spins up, is “online”, and can report its capacity. Using the STB Suite Original mode look at the device window or click the Scan System button to scan all of the storage controllers in your test system. Do you see the drive? Does it report a valid capacity? Does the drive information (manufacturer name, drive part number, firmware version) look reasonable?

If the drive does not show up at all on the test machine you must check all cabling and power for the drive test fixture. See if there are indicator LEDs on the drive that show any activity. Listen to see if you can hear that the drive is spinning, or gently feel if you can sense vibration from the drive.

If the drive does not report a reasonable capacity, for example if it reports it has zero blocks or a negative number of blocks than the drive may need a low-level format before continuing testing.

If the drive information is jumbled or wrong – for instance if the drive is a “SEAGATE” drive but the STB Suite is reporting it as “SEAGGGG” you may have a dead drive or you may have a cabling/termination problem. Try moving the drive to a different slot or connector or bus to see if the problem goes away.

Here is what you should see:



Can the drive tell us it has had problems?

Disk drives will store historical data which can be retrieved and analyzed. SCSI/FC/SAS drives store this type of information in LOG PAGES. ATA/SATA drives store this information in SATA SMART data.

SCSI/SAS/FC Log Pages

A quick way to see an overview of some of the more important Log Pages is to double-click on the drive in the device selection window to bring up the Device Information display. Select the **Error Data** tab and you will see historical data describing how much data has been read and written and the number of and type of errors that have happened during reads and writes –

Device Information

Host Adapter: 4 Target: 515 LUN: 0

Inquiry Information

Vendor: SEAGATE

Product: ST373455SS

Revision: S513

s/n: 3LQ1271W

Device Type: 0 - Direct Access

Removable: No

Multi-Ported: Dual Ported

Data Width: 8 bit

Taq Queing: Supported

SCSI Level: 5

SAS/wwn: 5000000004945CAC

Statistics Error Data EVPD Data Mode Pages

Write Error Counters

Total KB Written	1810597266	Errors Recovered w/o Delay	0
Errors Recovered w/ retries	0	Errors Recovered w/ ECC	0
Total Times Recovery Invoked	0		
Total Unrecovered Errors	0	Total Errors Recovered	0

Read Error Counters


Total KB Read	832120180	Errors Recovered w/o Delay	763558
Errors Recovered w/ retries	0	Errors Recovered w/ ECC	1
Total Times Recovery Invoked	763559		
Total Unrecovered Errors	0	Total Errors Recovered	763559

Clear Values Print

OK Print All Info


As a general rule – uncorrected errors are always a bad thing to have. Uncorrected errors usually will cause the LBA in question to be marked as bad.

For SCSI/SAS/FC drives this will mean an increase in the drives G defect list. On the Device Information display click the Statistics tab and note the number of G List defects –



As another general rule – good disk drives don't have any grown defects.

For a detailed “raw” data view of every log page a drive has you can right-click on the drive in the device selection window, then from the Quick Command list choose **View Log Pages** –



Be sure to use the Browse button to select a log page definition file – the file “default.dat” is usually fine for any disk drive. The available Log Pages are shown on the left of the display, double-clicking on a Log Page will display that pages parameters on the right.

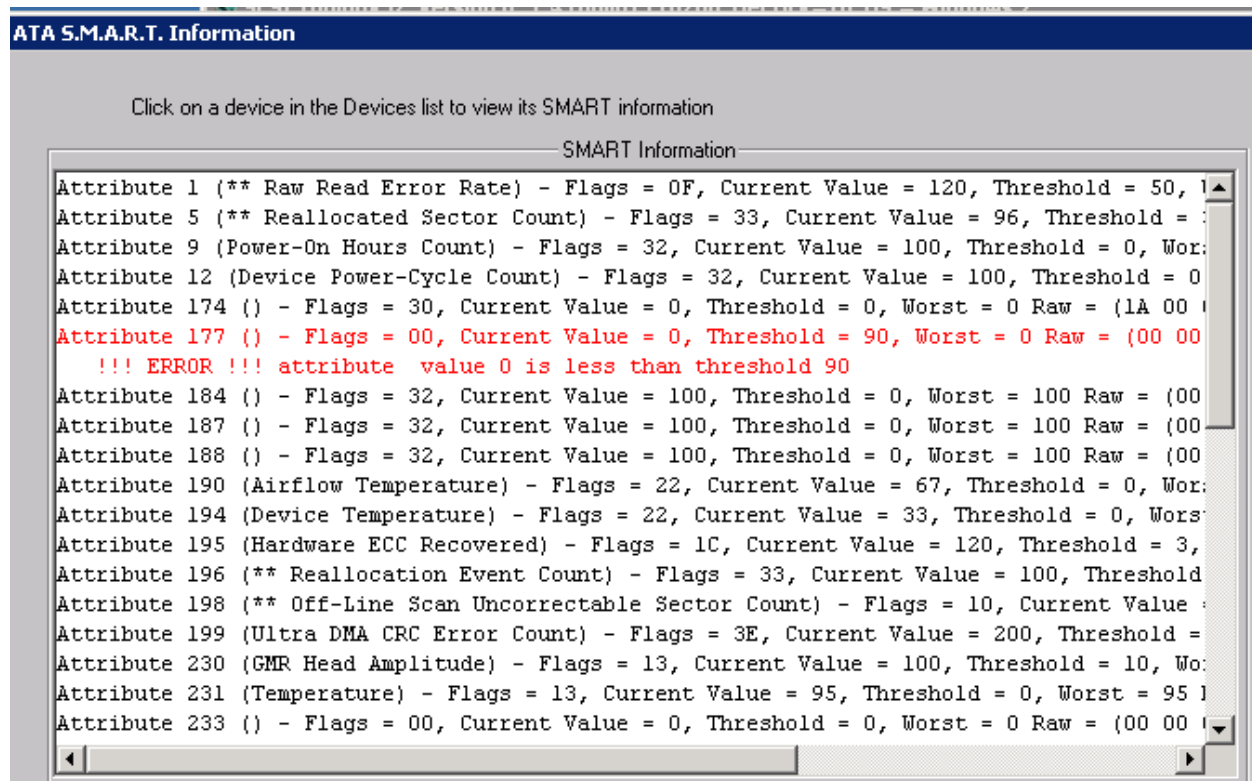
Note that you can save all of this information to a file. ***A good thing to do as you test drives is to build a database of the drives you’ve tested.***

ATA/SATA SMART Information

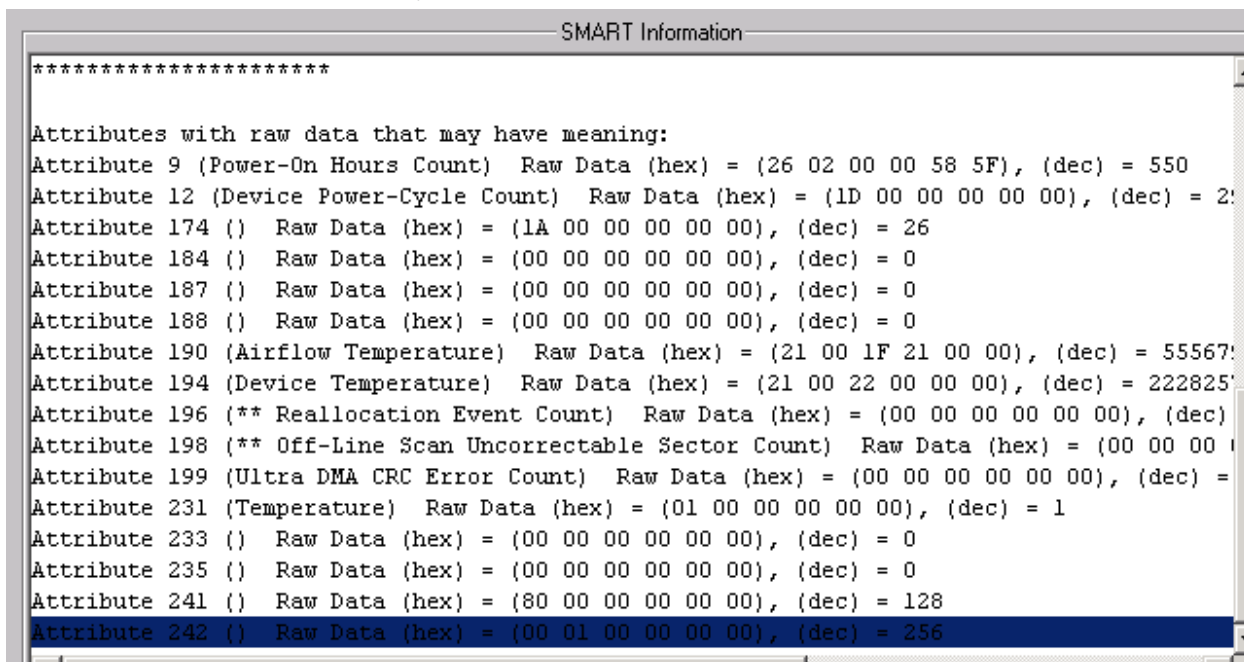
This same type of historical error data is found in the SMART data for ATA and SATA drives. To view and save this info go to the STB Suite main menu ATA/SATA->Commands->View SMART Data choice. Select the drive of interest from the lists to the right and the SMART data will be displayed on the left.

Note: to learn about how to interpret SATA SMART data look at <http://en.wikipedia.org/wiki/S.M.A.R.T.>

The top of the display will show all SMART attributes and will indicate pending problems



And at the bottom of the display you will see attributes which may indicate the actual number of errors or counts such as Power-On Hours, etc –



```
SMART Information
*****
Attributes with raw data that may have meaning:
Attribute 9 (Power-On Hours Count) Raw Data (hex) = (26 02 00 00 58 5F), (dec) = 550
Attribute 12 (Device Power-Cycle Count) Raw Data (hex) = (1D 00 00 00 00 00), (dec) = 29
Attribute 174 ( ) Raw Data (hex) = (1A 00 00 00 00 00), (dec) = 26
Attribute 184 ( ) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 187 ( ) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 188 ( ) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 190 (Airflow Temperature) Raw Data (hex) = (21 00 1F 21 00 00), (dec) = 55567
Attribute 194 (Device Temperature) Raw Data (hex) = (21 00 22 00 00 00), (dec) = 222825
Attribute 196 (** Reallocation Event Count) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 198 (** Off-Line Scan Uncorrectable Sector Count) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 199 (Ultra DMA CRC Error Count) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 231 (Temperature) Raw Data (hex) = (01 00 00 00 00 00), (dec) = 1
Attribute 233 ( ) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 235 ( ) Raw Data (hex) = (00 00 00 00 00 00), (dec) = 0
Attribute 241 ( ) Raw Data (hex) = (80 00 00 00 00 00), (dec) = 128
Attribute 242 ( ) Raw Data (hex) = (00 01 00 00 00 00), (dec) = 256
```

Can the drive reliably write and read data?

Obviously a disk drive must be able to reliably store and retrieve data. The best way to test these two functions is to run a test which first writes a known data pattern to every block on the drive. Then every block on the drive is read and using data compare each block of data is checked to insure that it is exactly the same as what was written.

Obviously writing and reading the entire drive will take some amount of time. Can you reliably determine if the drives write/read functionality is OK by checking less than the entire drive? Technically, probably not – what if there is a problem with a block which you didn't test? Statistically – maybe yes. The choice is up to you – balancing the accuracy of your results with the time it takes to complete a test.

The good news in this regard is the STB Suite Disk Manufacturing Module (DMM) is extremely efficient and fast – it can test many drives at once. DMM will tell you exactly how long a given test is going to take to complete so as you test more drives you will soon learn how many drives per hour or day you will be able to test to your company's specification.

Another choice to be made concerning write/read testing is the **access method**. The most basic access method is **sequential** – the test starts at LBA 0 and progresses sequentially through to the last block. Another access method is **Random**. As the name implies, random access moves through the drives blocks in a random manner. An advantage of random access is that it will generate more vibration in the drives under test, which will stress the drive harder. A new STB Suite access method is **CPAM**. CPAM is a

method which creates random access and also guarantees that every block on the drive will be accessed once and only once.

Getting started with DMM is covered in earlier articles and in videos on our web site at :

<http://www.scsitoolbox.com/Training/>

The STB Original mode also has a number of canned tests, more appropriate for testing a single drive at a time. Select a drive then click the top menu Disk-Tests choice pulldown to see a list of available tests – for example, the Quick QC test checks write/read functionality at the beginning, middle, and end of the drive. The Quick Drive Profile Test will show a good overview of the drive -

The screenshot shows a window titled "Quick Drive Profile" with a close button (X) in the top right corner. The window displays the following information:

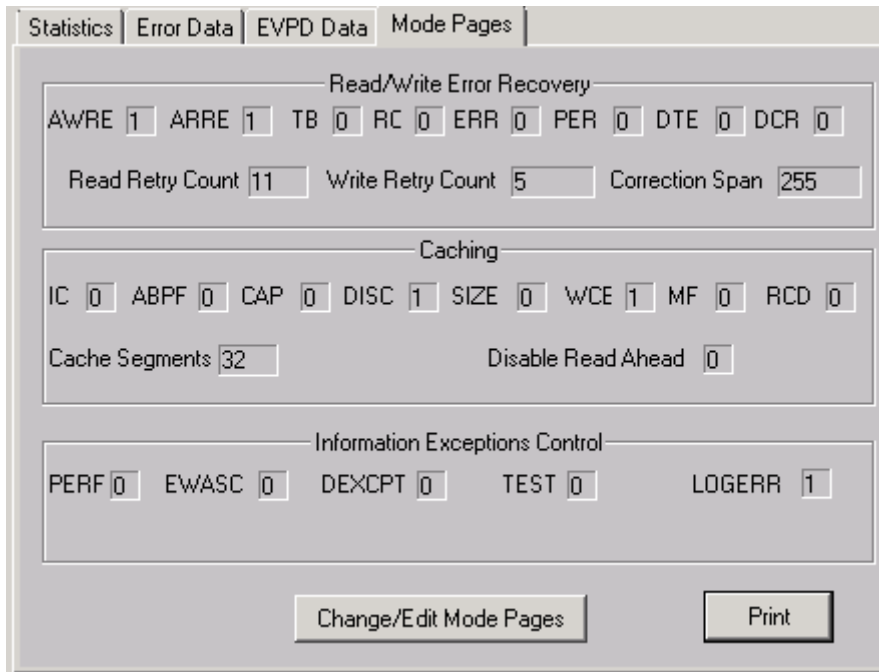
Vendor =	SEAGATE	Product =	ST373455SS	F/W =	S513								
Read Transfer Rate	Average =	86587.47	High =	153515.45	KB/S	Number of Blocks Read =	970						
Write Transfer Rate	Average =	183552.84	High =	191111.97	KB/S	Number of Blocks Written =	970						
Read Blocks	2 =	70.88	8 =	993.59	64 =	67124.85	128 =	148680.06	256 =	149139.97	512 =	153515.45	KB/S
Write Blocks	2 =	342.26	8 =	16201.22	64 =	125186.85	128 =	163836.77	256 =	183525.69	512 =	191111.97	KB/S
Average Seek	=	4.774893	Ms	Write Cache =	ON	Read Cache =	ON	Number of Segments =	32				
Number of Primary Defects =	1656	Number of Grown Defects =	0	Hard Errors =	0	Soft Errors =	0						
Status:	Finished												

Enable Destructive Tests

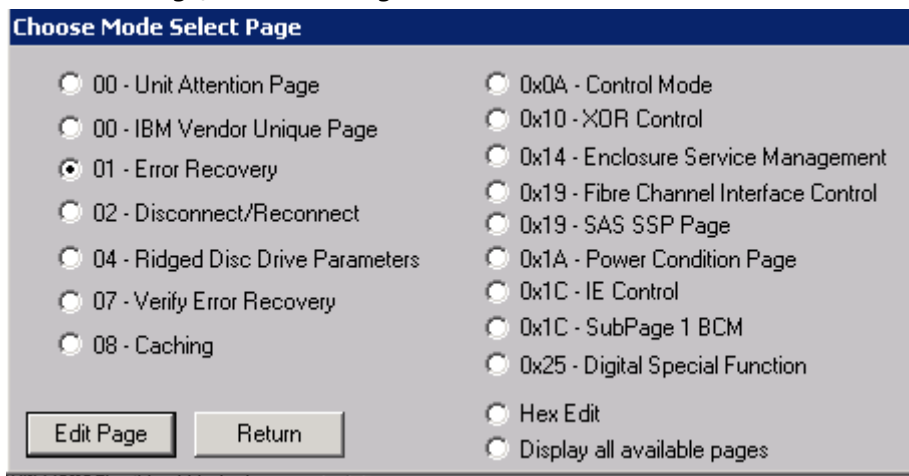
Exit Profile Drive File Output Cancel

Is the drive set up appropriately?

SCSI/SAS/FC drive behavior can be specified by Mode Page settings. To see the most common or important Mode Page settings go back to STB Original mode, double-click on your drive, and choose the **Mode Page tab** –



Settings such as enabling or disabling read ahead and write caching are shown, to change any settings click the **Change/Edit Mode Pages** button.



Note: in general you will want Write Caching (WCE) to be **ON**. This will greatly increase the write speed of the drive.

There are many settings available via Mode Pages. DMM has a feature whereby you can set up a golden drive with all mode pages set the way you or your customer defines – and then during DMM testing each drive under test will automatically have all of its Mode Pages set to match your golden drive.

What is the performance profile of your drive?


For a quick look at the performance of an individual drive you can select it in STB Original mode, then run a test and watch the real-time performance with Drive Watch – here is a view of the write

© Copyright 2011 SCSI Toolbox, LLC


Sales: 720.249.2641

General: 303.972.2072

performance of a drive:



DMM will reveal real-time performance metrics as well as logging them to log files. Here is a similar example to the above – a sequential write test :



Note that DMM tells you how long the current test step will take to complete – experiment with this and you will quickly get a feel for estimating how long any given test is going to take to complete.

And the DMM log file for this drive shows:

```
=====
>> SCSI Toolbox32, Version 8. 3. 4 (build) 110208, running on \\S2008x64-SSD <<
>> Default Driver: 10, Operating System: windows 7<<
=====

Test Date: 03/02/2011 11:23:07
Test Pass: 1

Device: 4:515:0 SEAGATE ST3734555S Serial: 3LQ1271W

results:

03/02/2011 11:23:07 TEST 1 of 1:
    Write Test; Sequential; for 1 Minutes
    Fixed-Length Transfers of 128 (0x0080) Blocks
    Start Block: 0
    Data Pattern: Decrementing
    Queue Depth = 1
    FUA = OFF

03/02/2011 11:24:07 Test Completed Successfully

    Transfer Rate: 124.32 MB/sec
    I/O Per Second: 1989.07 IO/sec
    Number of Blocks Transferred: 15276032
    Fastest Command Completion Time: 0.340 ms
    Slowest Command Completion Time: 18.760 ms
    Average Command Completion Time: 0.491 ms
    Standard Deviation of Command Completion Times: 0.227 ms

03/02/2011 11:24:07 PASSED
=====
```

Summary

Basic testing to confirm overall drive health is easily done with the STB Suite. Use Original mode to examine single drives at a time in detail/depth, and use DMM to test many drives at a time.

STB Original Mode Advantages:

- extreme depth of detail available to examine each and every drive setting
- versatile set of tests to get a quick snapshot view of drive health

STB Suite DMM Advantages:

- multi-threaded high speed multi drive testing
- can test multiple HBA's/controllers simultaneously
- extremely detailed logs generated for each drive under test
- easy to define any type of test sequence – test sequences can be saved and reloaded