



What's up with these numbers?

The Need for Performance Benchmarking Standardization



Esther Spanjer
Director, SSD Marketing



The Performance Landscape

- Bandwidth Performance Specifications
 - Sustained Sequential Read: Up to 250 MB/s
 - Sustained Sequential Write: Up to 70 MB/s
- Read and Write IOPS specifications (IOMeter Queue Depth 32)
 - Random 4 KB Reads: Up to 35 K IOPS
 - Random 4 KB Writes: Up to 3 K IOPS

Read/Write Mix?

Performance

Average Access Time	20-120 microseconds
Sustained Read Throughput	230 Mbytes/sec
Sustained Write Throughput	145 Mbytes/sec
Random IOPS Read Operations	45,000 IO/sec sustained
Random IOPS Write Operations	16,000 IO/sec sustained

Sustained?

IOPS?

Like other high-performance, enterprise-class flash drives, [redacted] drive is being positioned as an alternative to traditional 15,000-rpm serial-attached SCSI (SAS) hard disk drives. [redacted] said its new "performance-optimized" enterprise SSD can process IOPS (input/output per second) **more than 10 times faster** than the fastest SAS hard drive, with a sequential read rate of 230MB/sec. and a sequential write rate of 180MB/sec.

10 times faster than HDD?

Block Size?

Random?

Increased High-Speed Performance

- 800 Mbytes Full Duplex Burst Rate
- Up to 230 MB/sec Sustained Rate
- Up to 35,000 IOPS
- 30 to 100 µsec Access Time

Table 3: Sustained and Random Read/Write Performance

Parameter	Value
Burst Read	150 MBytes/sec ⁴
Burst Write	150 MBytes/sec
Sustained Read	Up to 115 MBytes/sec
Sustained Write	Up to 75 MBytes/sec
Random Read ⁵	6100 Input/Output Operations Per Second (IOPS)
Random Write	400 IOPS
Random 67% Read, 33% Write	1120 IOPS

Up to?


⁴ One megabyte, or MByte, equals 1,048,576 bytes.

⁵ Random performance values are based on 4 KByte transfers.

The 3 dimensions of SSD performance

3D IOPS Surface Profile


■ 0.0-500.0 ■ 500.0-1000.0 ■ 1000.0-1500.0 ■ 1500.0-2000.0
■ 2000.0-2500.0 ■ 2500.0-3000.0 ■ 3000.0-3500.0 ■ 3500.0-4000.0



- Performance depends on read/write mix **AND** block size

- 3D IOPS image can be created by SNIA SSSI Standard Reference Test Platform

Note: Shown 3D IOPS image courtesy of Calypso Systems






- Increased data fragmentation can result in performance drop over time
- Result of flash management algorithms that are not optimized for performing background operations
- Pre-conditioning is a must!

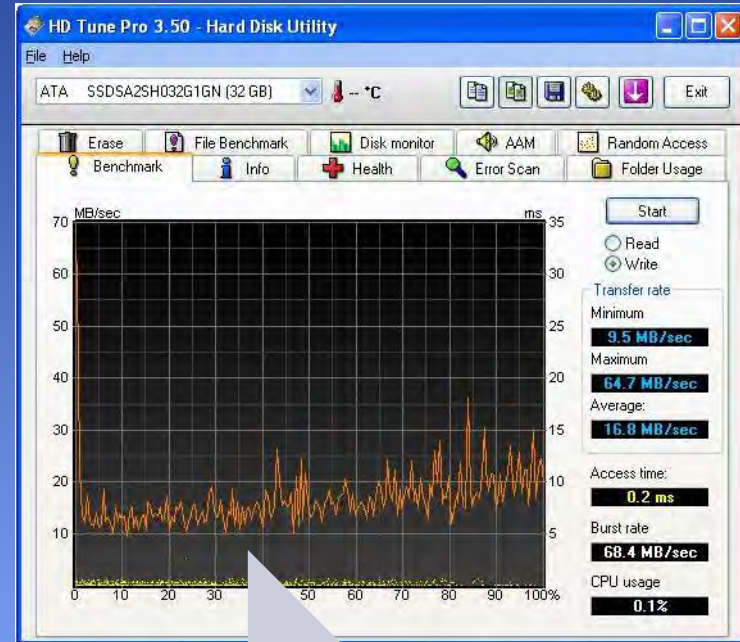
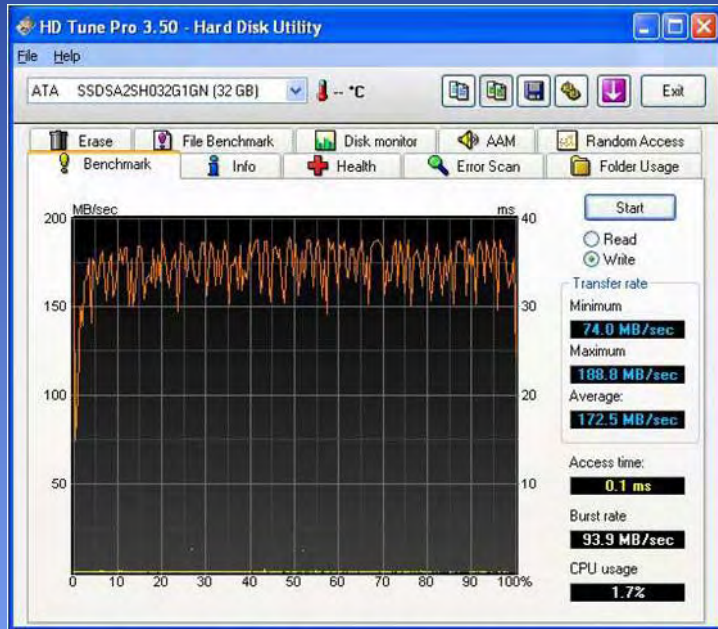


Not all F/W is equal

- Some flash management algorithms create non-deterministic latency
- Performance vs. time plot is important to characterize performance stability profile



Workload dependency - 1

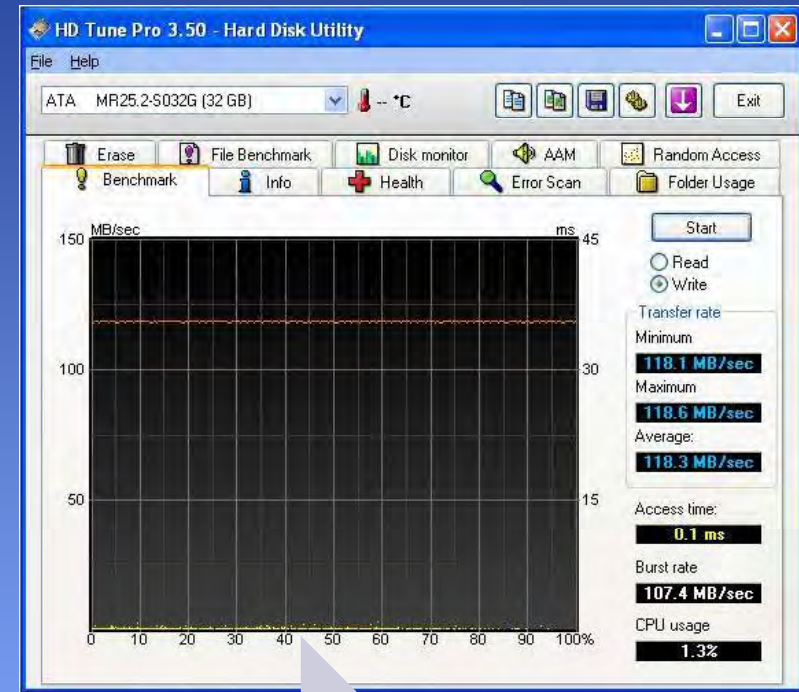
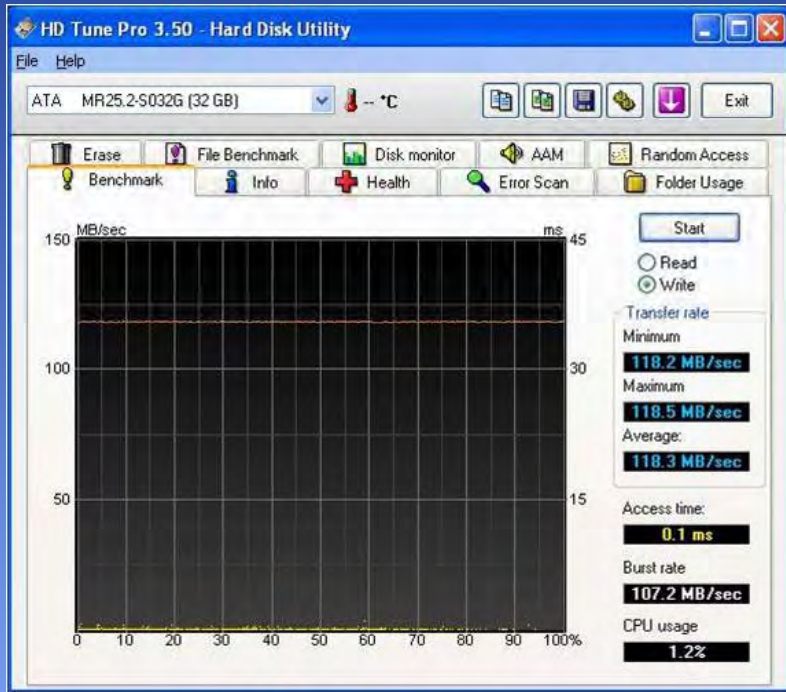


Step 1: HD Tune Pro
Sequential Write Test
Average write: 172.5MB/s

Step 2: IOMeter
Random Write Test

Step 3: HD Tune Pro
Sequential Write Test
Average write: 16.8MB/s
90% drop

Workload dependency - 2



Step 1: HD Tune Pro
Sequential Write Test
Average write: 118.3MB/s

Step 2: IOMeter
Random Write Test

Step 3: HD Tune Pro
Sequential Write Test
Average write: 113.3MB/s
4% drop

Benchmark Suites

	Test Suite	Commercial/ Client SSD	Enterprise SSD
PCMark	HDD Score, OS and application loading timing, user simulation (surfing web, windows, media player)	✓	
SysMark	System-level performance	✓	
IOMeter	Sequential workload server		✓
HDTach/ H2benchw	Performance stability, sequential performance, access time		✓
HD Tune	Performance stability, Sequential/Burst performance, Access Time		
Everest	Random Access Time (Read/Write)	✓	✓

Type of preconditioning and order of benchmarks can influence results

Performance benchmark steps

1. Pre-condition the drive
2. Run IOMeter for 3D IOPS view
 - Block size 512b-1MB
 - Entire Read/Write Mix range
3. Validate performance stability
 - Performance vs. time plot
4. Validate workload independency
 - Run sequential test, random test, sequential test
 - Run work load simulations (i.e. File server, web server. etc)
5. Run multiple composite benchmark tests to validate results

Standardization initiatives

- Technical Work Group (TWG) of SNIA
 - Performance Benchmark Standard: defines preconditioning, reference test platform, benchmark profiles, etc.
 - First draft available to public: 4Q09
- JEDEC 64.8
 - Specification for SSD endurance measurement
- SSDA
 - Testing of reliability (power cycling, data retention, endurance, etc) and OS compatibility (Windows 7)

Questions?