Configuring and setting up SSDs

Setup and Testing Guide

In this Guide

OCZ explores the issues and setup procedures for configuring and testing solid state drives (SSDs). This guide will detail known issues testing with existing benchmarks that are designed for traditional drives and how to properly gauge the real world performance benefits of the technology.

Known Issues:

The original release of Microsoft Vista has problems with properly using SSDs that are resolved with Service Pack1. Installation of SP1 results in a substantial performance boost in synthetic benchmarks and smoother operation of the system in standard usage.

Before using, please ensure that **Service Pack1 (Vista32 and 64)** is installed. Also, despite the fact that Vista contains generic chipset drivers, those drivers do not correctly enable any of the SATA controllers used on current motherboards, regardless of chipset manufacturer. Please make sure to install the latest **chipset drivers**, otherwise, especially the burst transfers will show a dramatic degradation and this will also affect sequential read and write performance

HDTune uses fixed strides across the array to measure small block sizes. If these blocks fall on page or chip boundaries, the time delays result in atypically low "calculated" performance spot data that are not representative of the drive's real performance but an artifact stemming from inadequate testing methodology.

HDTach's fixed strides between sample points causes similar issues as those outlined in HDTune, in this case, the periodicities of the stride pattern and the flash memory address space cause some wave-like interference patterns of the benchmark results that are also measuring artifacts rather than being indicative of performance.

AHCI: the Advanced Host Controller Interface allows Hot-Plugging and Native Commands Queuing as well as multithreaded access of the drive by applications. Enabling AHCI results in conflicts between the controller and the drive that are apparent as sluggish overall system performance.

Benchmarks that actually work are those that are not geared towards quick and dirty assessment of a HDDs performance based on optimization of the test algorithms to meet the typical HDD architecture. Examples are PCMark Vantage, Winbench 99 2.0 "Drive Inspection Test" or ATTO.

In Detail

Drive Setup:

As SSDs may require a slightly different setup that HDDs we recommend use of the latest BIOS and drivers. If you are having any real-life performance issue, please review our FAQ, most issues are easily resolved. Hard drive benchmarks as discussed below.

AHCI

AHCI is not official supported on OCZ SSDs and may under some circumstances affect performance, specifically during windows installation. Enabling AHCI can result in higher performance in synthetic benchmarks for SSDs and HDDs alike, but can cause hang-ups and intermittent freezes in SSDs since it allows multiple access requests to compete for a drive that is not made to address re-ordering of commands in the queue. We recommend AHCI is set to disabled in both Windows and in the BIOS. Native Command Queuing greatly increases the performance of standard rotational drives but it has no bearing on SSDs.

RAID performance with SSD

While RAID support for SSDs is not completely mature, the majority of chipsets using the most current drivers shouldn't have any issue. For maximum performance we recommend a stripe size 128kb or larger. A RAID Level0 x2 performance example is given below. Keep in mind that the Increased read and write speeds may be beneficial , but that the primary performance advantage of SSDs over HDD is in the seek time and IOPS.

Intel® Core™ 2 Duo E6300 1.86GHz, South Bridge: NVIDIA® 680i, 2GB DDR2 800MHz (5-5-5-12)

OS: Windows XP Pro SP2 M/B: EVGA

ATTO Disk Bench		
Model P/N	**READ (MB/s)**	**WRITE (MB/s)**
OCZSSD2-1C32G	**144.839**	**92.044**
OCZSSD2-1C64G	**132.888**	**89.587**
OCZSSD2-1C128G	**120.41**	**86.642**
RAID0 (x2)		
Model P/N	**READ (MB/s)**	**WRITE (MB/s)**
OCZSSD2-1C64G	**229.712**	**166.523**
OCZSSD2-1C128G	**240.861**	**174.416**

SSD vs. HDD:

The primary performance difference in SSD versus hard drives is related to seek and platter rotation delay times.

Hard Disk Drives are electromechanical devices that are using magnetic platters to store data within tracks and sectors. Access of tracks incurs seek latencies and access of the correct sector causes rotational latencies that in combination define the random access latency of a drive. SSDs do not use mechanical parts like conventional hard disc drives, and, therefore do not incur rotational and seek latencies. Likewise, track-switching delays that reduce the sustained read/write performance of an HDD don't exist. Unlike in the case of HDDs where sequential performance degrades from the outer tracks of any platter to the inner tracks, SSDs offer the same performance throughout the entire array of NAND Flash chips.

Conventional Hard Disk Drive Media Performance:

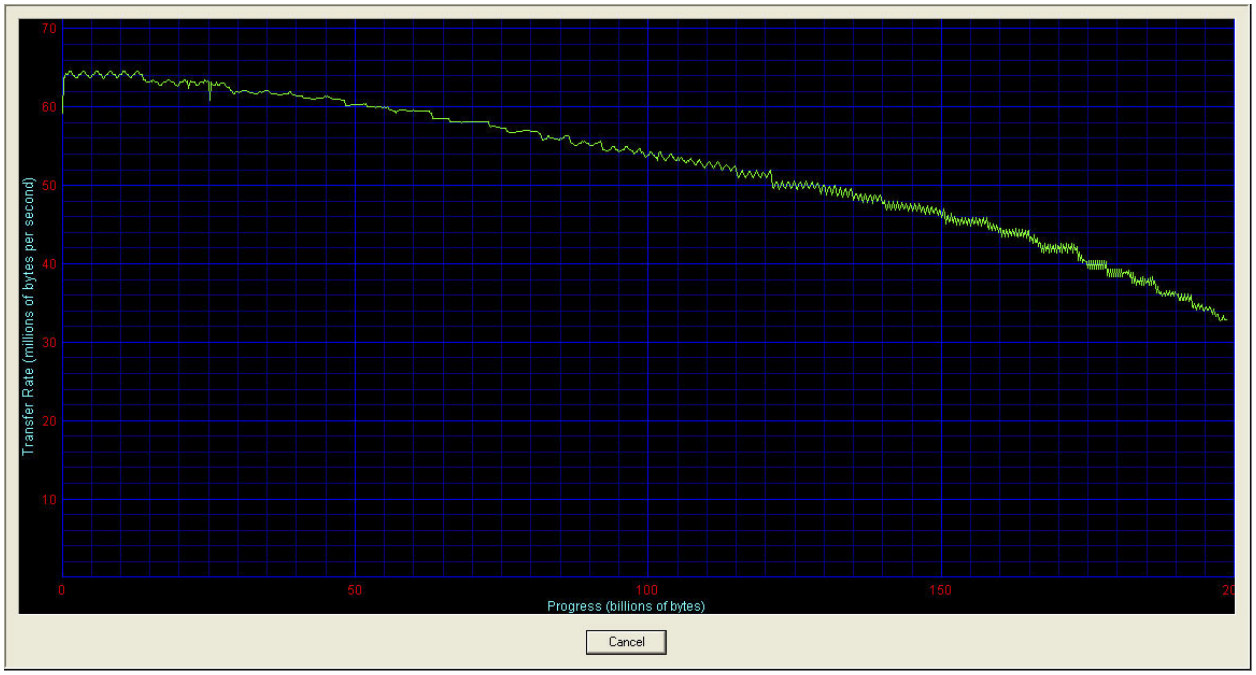

Typical sequential performance profile for a standard hard disk (7200 rpm - 200GB – 2.5” Hitachi Travelstar): The high performance seen at the outer diameter of the platters degrades to approximately 50% at the inner diameter. Because of this, all conventional HDD preferentially use the outer tracks and only use inner tracks when they start filling up. Because of the seek and rotational latencies defragmentation of HDDs is the best way to ensure continuously good performance.

Solid State Drive Media Performance

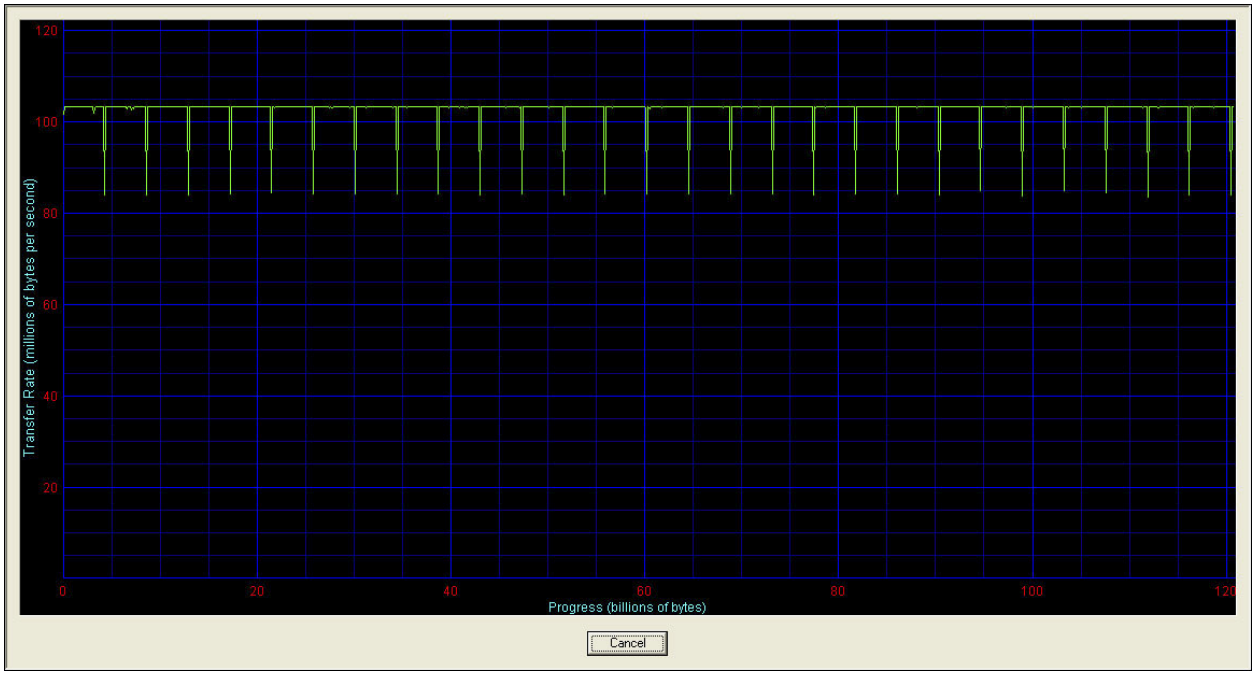

Typical sequential performance profile for a Solid State Drive (MLC NAND Flash-based 128GB – 2.5" OCZ Core SSD): No performance differences exist across the entire array. Therefore, SSDs have no preferred pattern of filling up. In fact, clustering data as it is done by defragmenting can cause problems because defragmenting can cause hot-spots in the memory array and negate wear-leveling, a measure used in SSDs to increase life span.

Differences in Addressing Compared to HDD:

Conventional hard disk drives utilize rotating platters that are divided into cylinders (tracks), where each track contains a certain number of sectors. Each sector is 512 Bytes long. In addition, HDD platters contain servo bursts that are reserved areas for track information etc. Most HDD benchmarks are configured such that the benchmark effectively tests the data area, if a test block is given that spans across a servo burst or else from one track to the next, then the benchmark usually generates inconsistent benchmark results.

SSDs use a completely different address system, which is based on "chips," "planes," "blocks," and "pages." Each page contains 4 Kbytes of data, however, because of the parallelism at the back end of the controller, every access includes simultaneous opening of 16 pages for a total accessible data contingent of 64 Kbytes. Each block contains 128 pages and the number of blocks per chip defines the total density of the chip.

Benchmarking SSD Drives:

Before benchmarking, please ensure that **Service Pack1 (Vista32 and 64)** is installed. Also, despite the fact that Vista contains generic chipset drivers, those drivers do not correctly enable any of the SATA controllers used on current motherboards, regardless of chipset manufacturer. Please make sure to install the latest **chipset drivers**, otherwise, especially the burst transfers will show a dramatic degradation and this will also affect sequential read and write performance.

Benchmarking SSDs using tools developed for HDDs causes several unique problems that need to be solved by developing benchmarking software that catches up with the technology.

As mentioned, SSDs use different strategies and data geometry than conventional HDDs. This causes some functional differences and, more importantly, makes some benchmarks inadequate, particularly those that were optimized for the standard platter configuration of HDDs. Due to these addressing issues, some benchmarks could show radically different results also on the transfer graphs and or average the performance values incorrectly.

Needless to say that the same algorithms applied to a functionally totally different device will not render the same "realistic" performance values, on the contrary, many of the test points will fall within one block but others will span from the end of one block to the beginning of another block. That will cause delays in the completion of the reads or writes and since the test samples are relatively small in size, it will result in low "calculated" performance values. Because the stride size is constant in most benchmarks and the page size is constant, too, these values will result in a sawtooth pattern of the performance graph, simply as a consequence of the periodicity of the two address patterns. This may be less apparent on SRAM cache equipped SSDs but nonetheless gives "artificially" low values.

Some benchmarks appear to use test patterns that don't seem to work well with SSDs and, thus, generate artifacts. For example, HDTune consistently generates burst speed numbers that are below the sustained transfer rate values. From a technical point of view, this is not possible since the measuring of internal media performance (sequential transfers) is limited by the bus interface (burst rate) bottleneck. This type of numbers needs to be viewed as artifact and the benchmark cannot be viewed as representative of the SSD performance.

We have generated some basic data on benchmarks

Benchmark notes:

PC mark Vantage

The PC mark vantage HDD test is intended to mirror real life applications and is recommended as a key comparator versus other SSDs and HDD. This is currently the benchmark of Choice for SSDS or HDDs.

Standard HDD benchmarking software

The industry standard benchmarks for the last decade have been Winbench99, especially "Drive Inspection Test," IOMeter, and ATTO. For general media performance WB99 2.0 is still the test of choice since instead of spot checking, the entire storage media is tested. This causes the test to run longer than HDTach or HDTune but the results are accurate representations of the media performance across the *drive instead of extrapolations of spot data.*

ATTO

ATTO is the preferred Standard benchmarking software as while it runs spot data is has several advantages over other HDD benchmarks including the fact that It shows the specific speed of each file size transfer and shows a true average , this benchmarking product also doesn't seem to favor SRAM over DRAM caches .

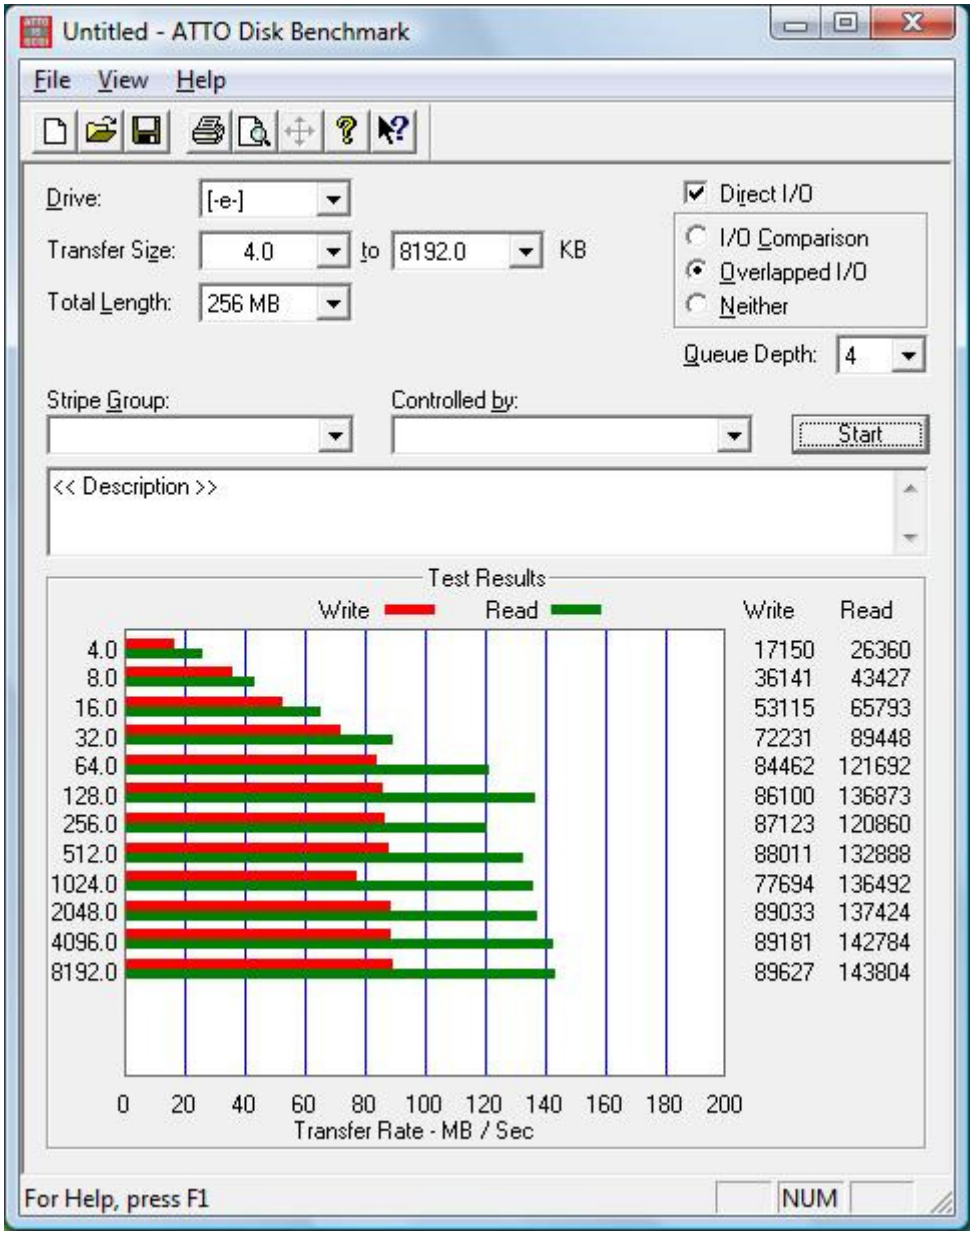

OCZ Core 64GB SSD Performance measured by ATTO (single drive)

HD bench

HD bench allows user choice of the file size and while simple, provides extremely accurate results

HDTACH and HD Tune

Addressing issues as described above me into play with these benchmarks making them basically unusable for most SSDs. While this test can be run quickly, they will tend to over or under report the true R/W speeds. However they are useful for CPU utilization and Seek time measurements. This is compounded by the fact that they use periodic sampling to extrapolate speed. While on some chipsets the periodic Sampling may show less artificial degradation on expensive SRAM CACHE enabled drives, but nonetheless skews results in favor of rotating media by using algorithms that optimize the sampling

periodicity according to the platter geometry. This makes HDTach and HDTune very accurate for HDD testing but inadequate for SSD tests.

WB99 2.0 with Core (64 GB) single drive

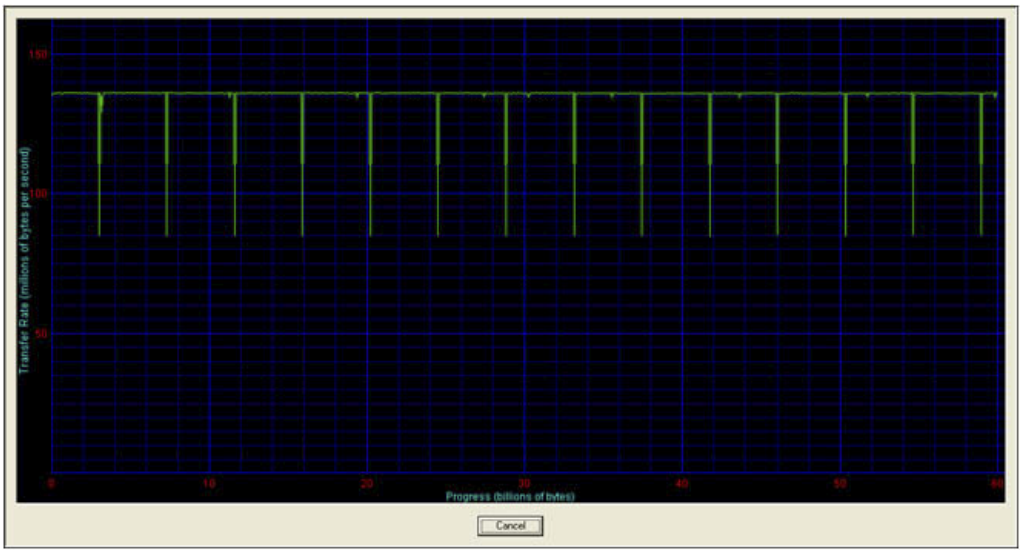

The same SSD measured with HDTach: Note how three of the chip-switch intervals that are apparent in WB99 are hit by HDTach's periodic sampling.

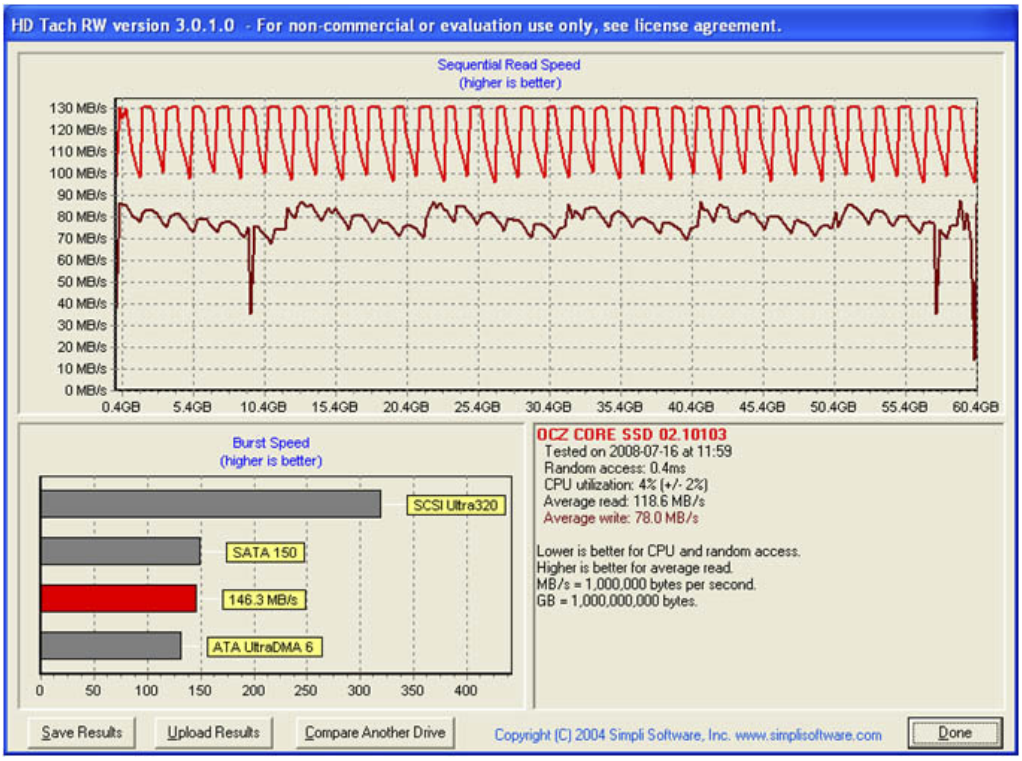

Also note that the graph shows a "smoothed" curve, using data point interpolation of small samples to generate the read and write speed graphs. If the same test is repeated in HDTach using the "Full Test" that substitutes variable length test blocks across the entire address range for the fixed test spot-

checking and data point interpolation, the picture changes dramatically to look like WB99. However, in this case, a median of the different transfer rates at different test block sizes is used. As a result, the average transfer numbers are lower than what the drive is actually capable of.

SiSoft Sandra File System Test

SiSoft Sandra takes a long time to run but provides consistent results using a 2 MB test sample size.

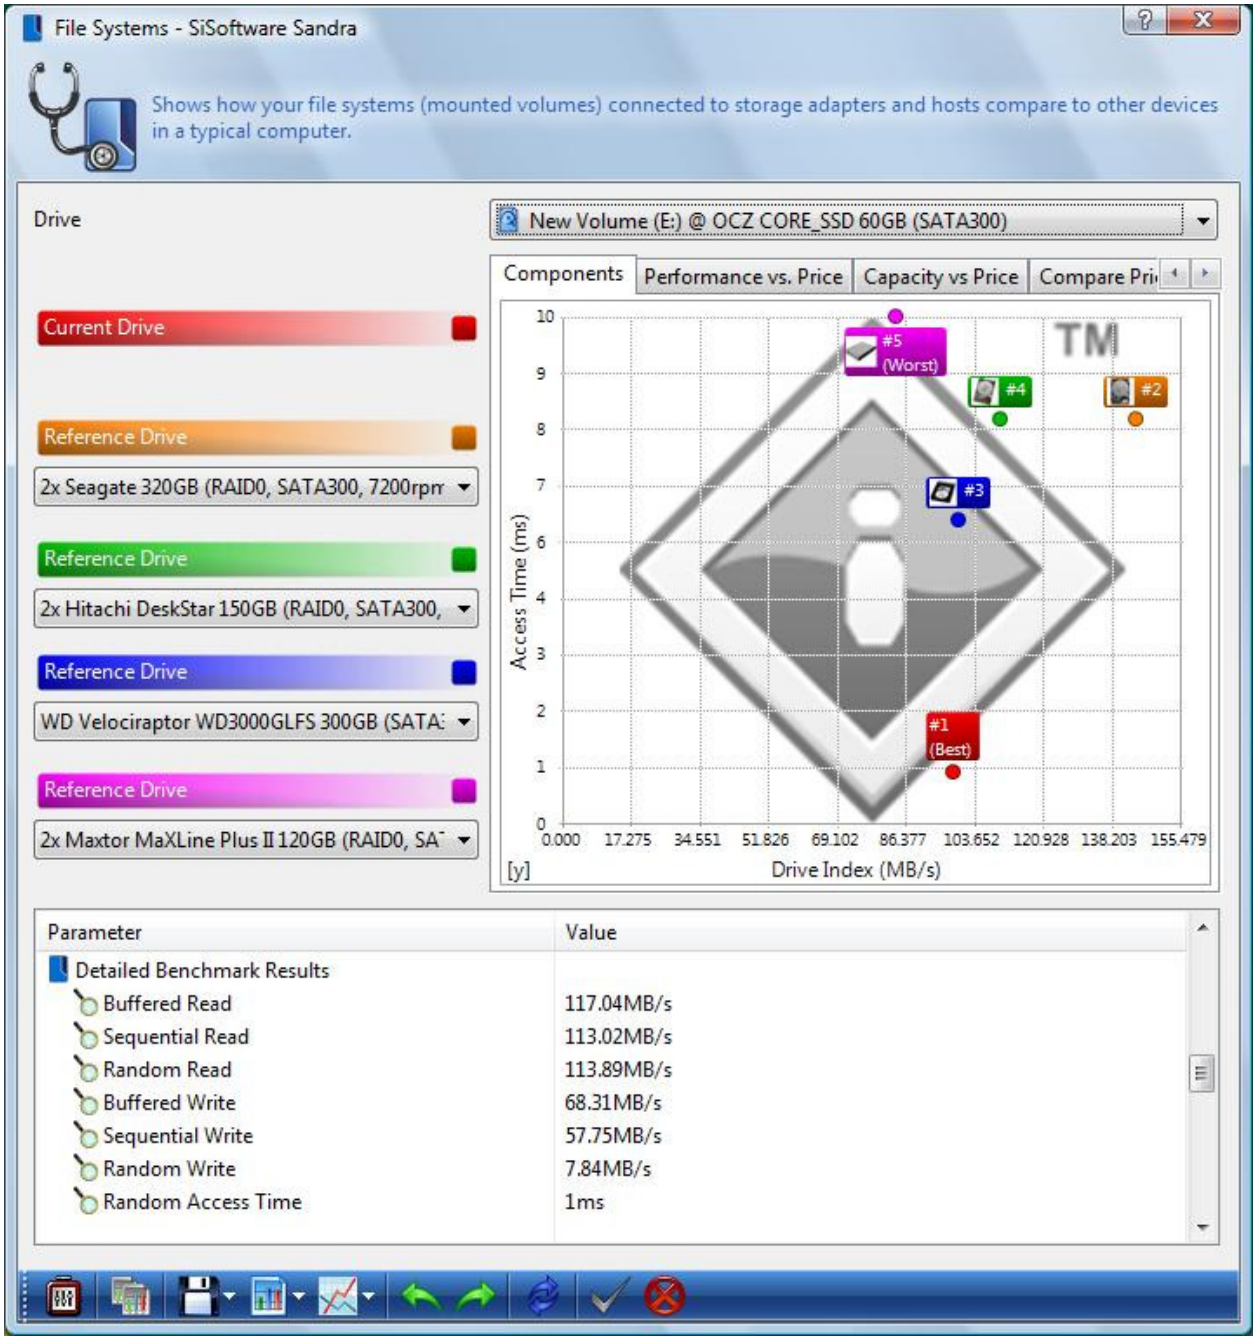

Note that the numbers for SiSoft are again different from those in HDTach or WB99.

Conclusion:

Every benchmark uses different methods of measuring performance and some benchmarks appear inadequate for measuring SSD performance, others are somewhat better but there are no benchmarks out there that are specifically geared towards SSDs.