# Comparison of Drive Technologies for High-Transaction Databases

August 2007

By Solid Data Systems, Inc.

**Contents:**



SolidData®

## ABSTRACT

Database transactions are by nature random and usually come with small block sizes. Cost and performance tradeoffs between hard drives, flash-based Solid-State Disks (SSDs), and DRAM-based SSDs are worth investigating. In this paper, we discuss considerations in choosing the right technology and provide one viewpoint on performance comparison using industry-standard benchmarks and handy graphical guides for estimating read and write Database performance improvements.

High-transaction databases are typically comprised of small records (i.e., 4 or 8 Kbytes) that are often accessed randomly. Because the records are brief when compared to the time required to reach the data location, mechanical disk drives are paced by their ability to locate and retrieve information on the disks (i.e., disk access time). Disk access time becomes the dominant reason for slow database application performance, often causing the CPU to wait for disk I/O to complete.

The problem is compounded by ever-increasing disk drive capacities because, as companies take advantage of larger drives, they are accessing fewer spindles for the same amount of data. With a 50 GB drive, less than 0.004% of the drive is accessible each second when randomly accessed for database-sized records. Striping information over multiple drives can increase the amount of data that is accessed, but transaction integrity requirements of the database often conflict with broad, simultaneous access.

Solid State Disks (SSDs) are an excellent solution to the problem, and provide various levels of performance based on the underlying memory technology.

## COMPARING DRIVE TECHNOLOGIES

### Hard Disks

Table 1 indicates the random performance that can be expected from a single rotating disk. When block sizes are 4 or 8 kB – the typical database size – these disks can read and write data at only 1 to 2 MBps. This is because the seek time (the time required for the heads to move mechanically from track to track), plus the rotational latency (the time for the block of interest to rotate under the read/write head), is 5 to 10 milliseconds for modern disk drives. Note as well that if the block is as large as 128K, then the rotating disk performance improves. Thus, rotating disks work well when reading or writing large files, but are limited when used for database-sized blocks. Needless to say, rotating disks are very economical for their capacity, but limited in performance.

| RANDOM READ BENCHMARK | | | | RANDOM WRITE BENCHMARK | | | |
|---|---|---|---|---|---|---|---|
| Block Size | Read IO/s | Read MB/s | Avg Service Time - ms | Block Size | Write IO/s | Write MB/s | Avg Service Time - ms |
| 512B | 185 | 0.09 | 10.4 | 512B | 290 | 0.14 | 6.7 |
| 1K | 185 | 0.18 | 10.5 | 1K | 290 | 0.29 | 6.5 |
| 2K | 182 | 0.37 | 10.5 | 2K | 283 | 0.57 | 6.9 |
| 4K | 175 | 0.70 | 10.8 | 4K | 280 | 1.12 | 6.3 |
| 8K | 176 | 1.41 | 10.9 | 8K | 284 | 2.27 | 6.2 |
| 16K | 172 | 2.75 | 11.0 | 16K | 264 | 4.23 | 6.3 |
| 32K | 170 | 5.44 | 11.0 | 32K | 237 | 7.58 | 6.6 |
| 64K | 152 | 9.76 | 11.0 | 64K | 211 | 13.51 | 6,5 |
| 128K | 132 | 16.96 | 11.2 | 128K | 183 | 23.48 | 8.0 |

**Table 1**

**Flash-Based SSDs**

Because they have no moving parts, flash-based SSDs are a popular option in ruggedized applications or under high *g* load as in airborne applications. On reads, their performance falls between a conventional rotating drive and a DRAM-based SSD. On writes, however, they are generally the slowest of any disk technology at any block size. Performance figures are shown in Table 2.

| RANDOM READ BENCHMARK | | | | RANDOM WRITE BENCHMARK | | | |
|---|---|---|---|---|---|---|---|
| Block Size | Read IO/s | Read MB/s | Avg Service Time - ms | Block Size | Write IO/s | Write MB/s | Avg Service Time - ms |
| 512B | 1315 | 0.66 | 1.4 | 512B | 22 | 0.01 | 92.5 |
| 1K | 1217 | 1.22 | 1.5 | 1K | 22 | 0.02 | 91.7 |
| 2K | 1206 | 2.41 | 1.5 | 2K | 21 | 0.04 | 92.3 |
| 4K | 1075 | 4.30 | 1.7 | 4K | 21 | 0.09 | 94.5 |
| 8K | 906 | 7.26 | 2.0 | 8K | 21 | 0.17 | 92.5 |
| 16K | 666 | 10.66 | 2.8 | 16K | 21 | 0.34 | 93.7 |
| 32K | 447 | 14.33 | 4.2 | 32K | 21 | 0.68 | 102.1 |
| 64K | 322 | 20.62 | 5.9 | 64K | 19 | 1.23 | 106.7 |
| 128K | 204 | 26.16 | 9.5 | 128K | 18 | 2.37 | 113.2 |

**Table 2**

**DRAM-Based SSDs**

As shown in Table 3, DRAM-based SSDs are the performance champions for database use, producing very high quantities of reads and writes per second and very low service times. Often, large databases may be performance limited by a few hot files or a heavily used temporary sort area. Placing just these files on DRAM-based SSDs can provide substantial overall performance gains, sometimes as much as an order of magnitude. Note as well that as the block size increases, the benefit over conventional rotating drive diminishes. Thus, DRAM-based SSDs are typically used to improve random small block performance. They are the most expensive of all drives due to the use of internal uninterruptible power supplies and hard drive backup. However, DRAM-based SSDs are still a reasonable solution in many cases because only a small percentage of the database must be stored on them.

| RANDOM READ BENCHMARK | | | | RANDOM WRITE BENCHMARK | | | |
|---|---|---|---|---|---|---|---|
| Block Size | Read IO/s | Read MB/s | Avg Service Time - ms | Block Size | Write IO/s | Write MB/s | Avg Service Time - ms |
| 512B | 7388 | 3.69 | 0.2 | 512B | 7238 | 3.62 | 0.2 |
| 1K | 6794 | 6.79 | 0.2 | 1K | 6015 | 6.01 | 0.2 |
| 2K | 5752 | 11.50 | 0.2 | 2K | 5353 | 10.71 | 0.3 |
| 4K | 4091 | 16.36 | 0.4 | 4K | 4184 | 16.74 | 0.4 |
| 8K | 2959 | 23.68 | 0.6 | 8K | 2875 | 23.00 | 0.6 |
| 16K | 1771 | 28.35 | 1.0 | 16K | 1773 | 28.37 | 1.0 |
| 32K | 999 | 31.98 | 1.8 | 32K | 1004 | 32.14 | 1.8 |
| 64K | 531 | 34.01 | 3.5 | 64K | 540 | 34.57 | 3.5 |
| 128K | 277 | 35.47 | 6.9 | 128K | 281 | 36.08 | 6.8 |

**Table 3**

## WHEN SPEED COUNTS

While less expensive than DRAM-based SSDs at an equivalent capacity, hard drives or flash-based SSDs may not deliver enough overall performance in some applications.  The leading choice for overall performance in transactional real-time database environments is DRAM-based SSD.

## ABOUT THE AUTHOR

**Wade Tuma** is the founder and CEO of Solid Data Systems, and is a technical visionary in the SSD industry. Prior to launching Solid Data Systems, Wade was the founder, president, and CEO of Compower Corporation. He has more than three decades of experience including significant work in very high-speed digital logic, radio frequency, video, and color-generation systems for National Semiconductor, Atari, and Apple.  He earned a BS degree in Electrical Engineering from Cleveland State University.

To learn more, contact Wade at wtuma@soliddata.com or visit www.soliddata.com.

## APPENDIX

The following tables provide a graphical comparison of the three drive technologies for reads and writes.

READS

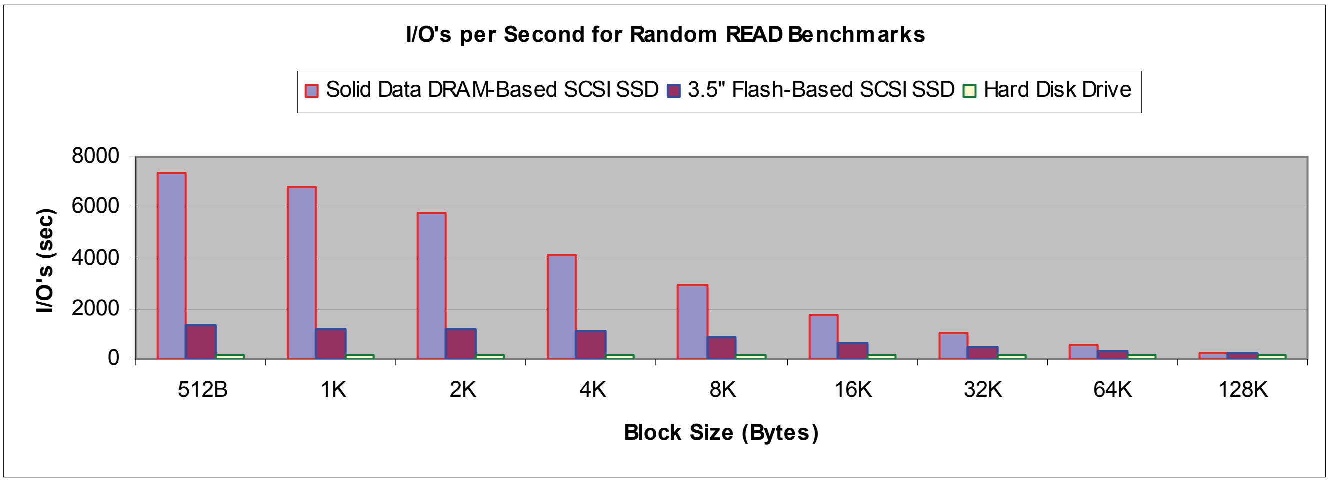


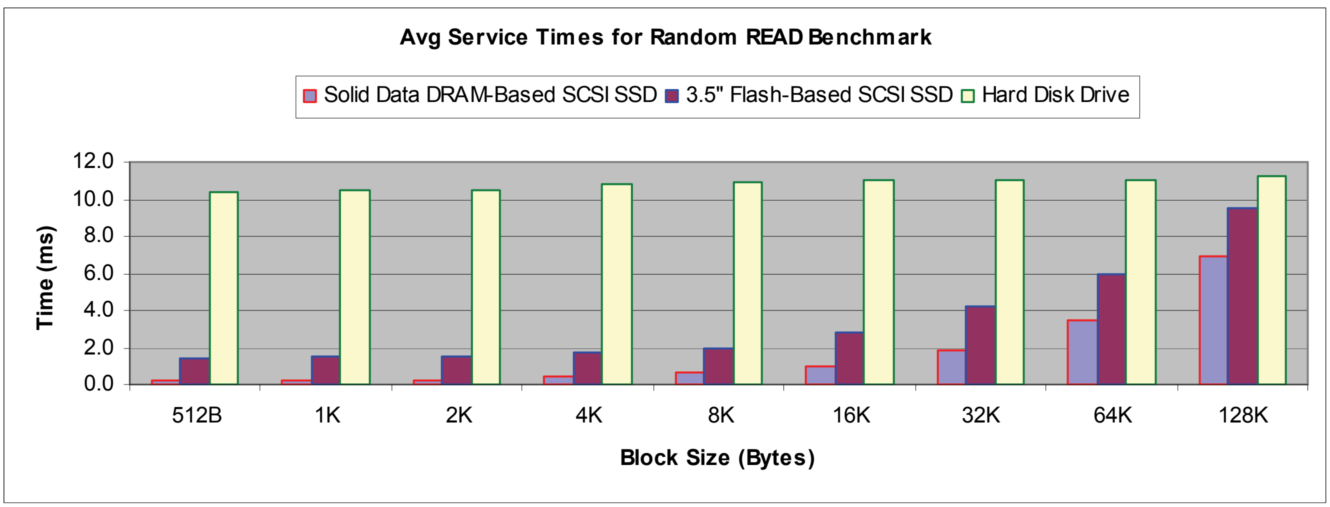


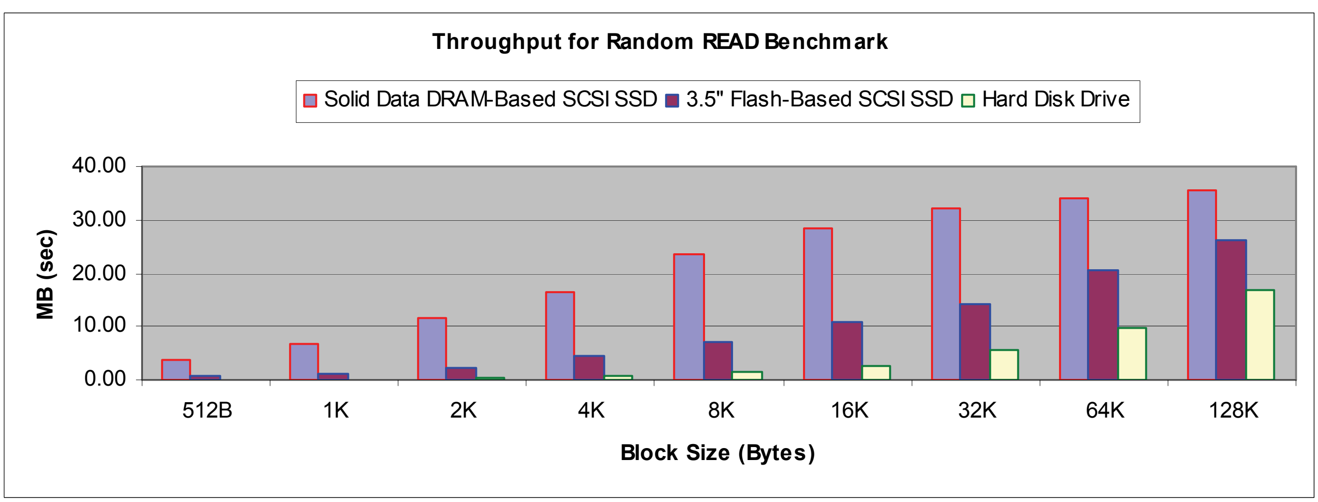

# APPENDIX (CONT'D)

## WRITES