Optimizing Solid State Storage With *HyperFast™ Technology*

The Problem

Solid State Drives (SSD) have several issues that come into play when used as the primary storage medium in place of traditional magnetic Hard Disk Drives (HDD); specifically regarding the operating system's logical file system. To date, file systems have been designed under the pretext of HDDs, and do not properly account for new NAND Flash / SSD technology.

This leads to issues that negatively impact mass market adoption of NAND Flash memory into the consumer, small/medium businesses, and enterprise markets where legacy file systems, designed for HDDs, will dominate the landscape for many years.

Problem Analysis

The principle issue is write speed degradation due to free space fragmentation. Small free spaces scattered throughout a volume at the logical level cause the file system to write a file in fragments to those small free spaces. ***This will degrade write performance as much as 80% to that solid state storage device*.** Fragmented free space forces the writing of blocks that are not advantageous to write or erase block boundaries' of SSDs, even if the application software would otherwise use large buffers.

Diskeeper Corporation has tested current SSDs, from a wide cross section of manufacturers for file write speed degradation as a function of free space fragmentation. NTFS, the most common Windows file system by far, fragments free space rather aggressively over a few months and then grows slowly thereafter when not maintained (figure 1).

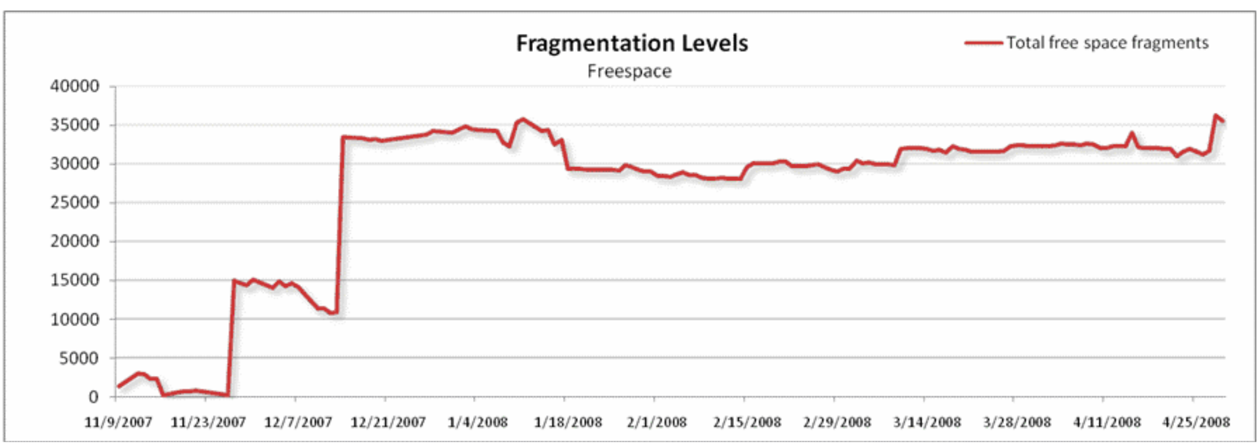

Figure 1: Accumulation of free space fragmentation through normal system use.

Write performance decreases proportionately as free space fragmentation increases. The graph below (Figure 2) shows the resultant 80% write speed reduction on a variety of drives due to this effect. The test included copying a 1GB file on to an SSD with fragmented free space. The graph axis depict the I/O write throughput in megabytes per second (MB/s) in relation to the number of file fragments the file is forced to be written in due to non-contiguous free space clusters. While a brand new SSD device may offer write performance in the 80 MB/s range to start, after a few weeks of use, performance will quickly deteriorate to 35 MB/s. Over the span of a few months, write speed becomes painstakingly slow at an abysmal 10 MB /s.

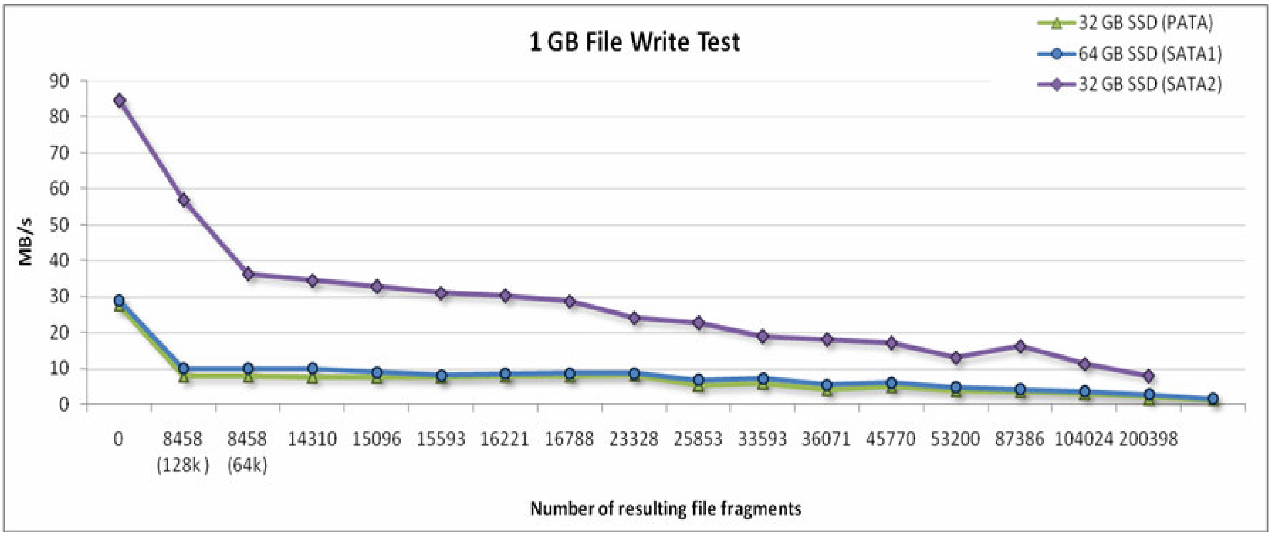

Figure 2: Write I/O throughput degredation due to fragmented free space.

It is important to note that the performance loss due to fragmented free space affects most every NAND Flash memory drive shipping today.

The Software Solution: HyperFast

As described, the issue originates with how an operating system interacts with the storage hardware. The nature of modern operating systems, their regular updates, and the applications that run on them, do so with no forethought to the effect of free space fragmentation and its unique impact on SSD NAND flash. Empirical evidence proves that, NAND flash, while moderately immune to effects on read-based file fragmentation, is extremely susceptible to write speed degradation when the free space is moderately to heavily fragmented.

Using Windows FSCTLs (file system controls), originally co-written by Diskeeper Corporation and Microsoft, HyperFast delivers 100% safe automatic maintenance of the file system, keeping a low level of free space fragmentation through specific optimization techniques which preemptively force the file system to write sequentially rather than randomly.

<u>Performance:</u>

This technology dramatically improves SSD performance. Specific tests were performed to demonstrate the improvement. The test depicts an environment with free space fragmentation in a range as might be expected from about 6 months of typical system use.

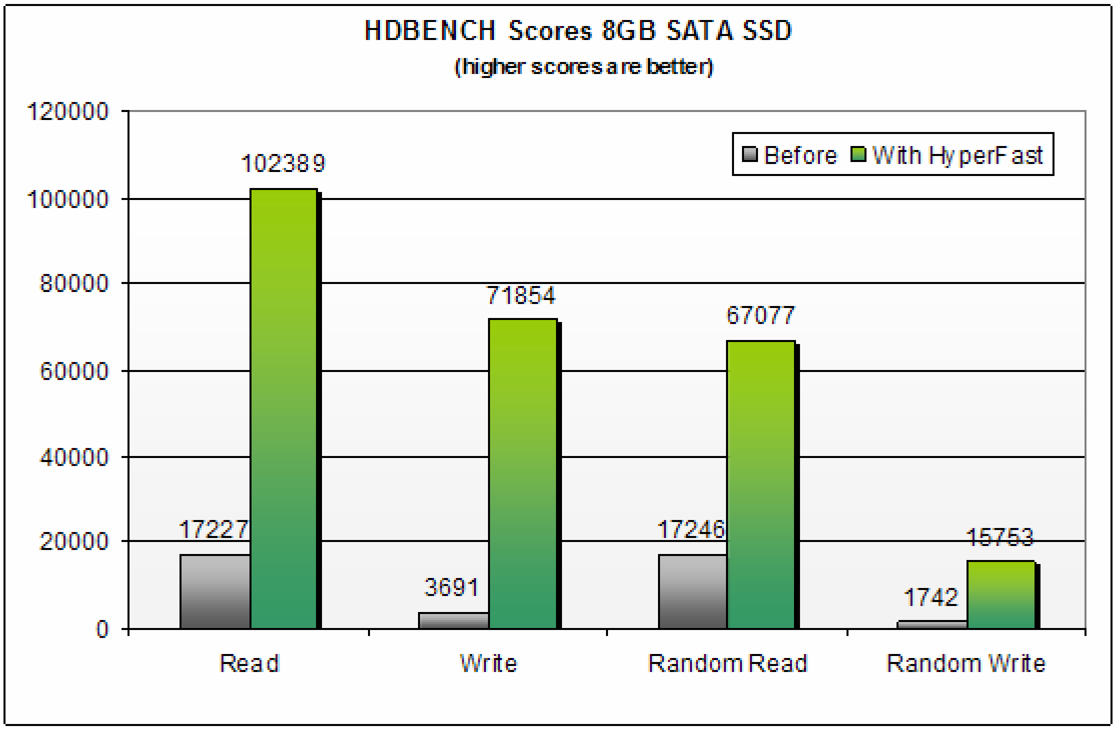

Figure 3: Benchmark scores on 8GB SATA SSD with HDBench™ software.

The results on an 8GB SSD (SATA) device record performance improvements of 5.9x faster reads, 19.5 faster writes, 3.9 faster random reads, and 9.0 faster random writes.

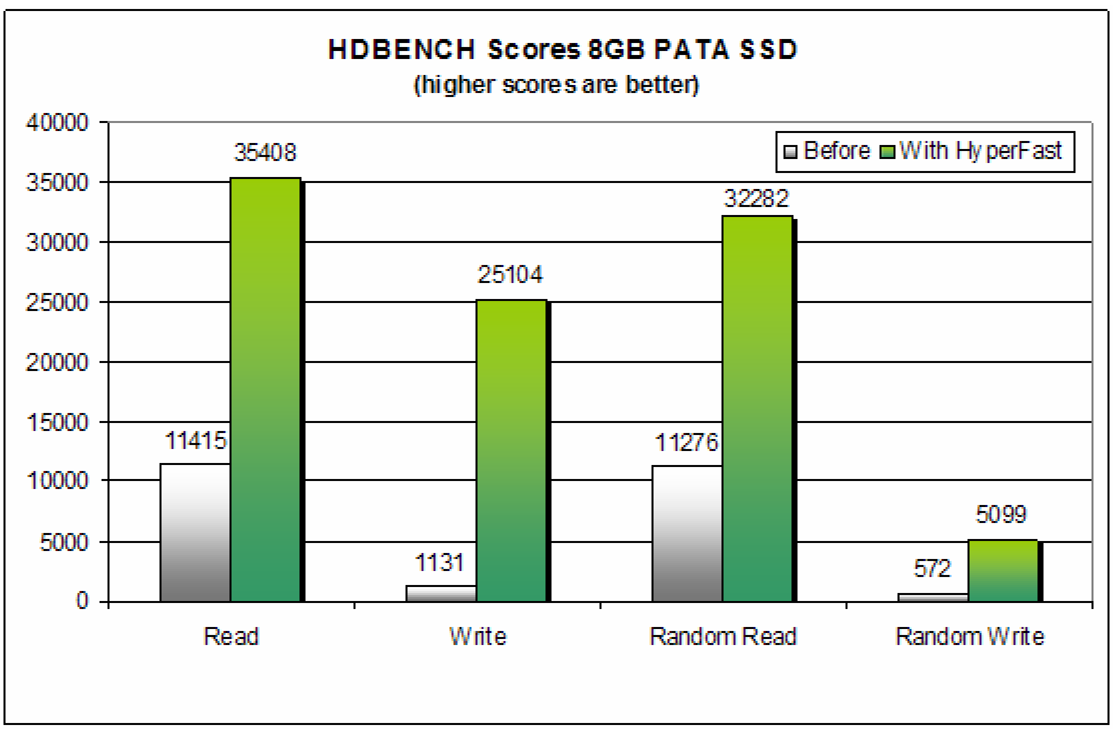

Figure 4: Benchmark scores on 8GB PATA SSD with HDBench software.

The results on an 8GB SSD (PATA) device record performance improvements of 3.1x faster reads, 22.1 faster writes, 2.8 faster random reads, and 8.9 faster random writes.

With SSDs typically offering sequential write I/O throughput (MB/s) far superior to their random write I/O, HyperFast's actions to proactively consolidate fragmented space in the file system improves performance by preventing sequential write I/O from splitting into what would then otherwise be passed to the disk subsystem as random write I/O.

Boot Up:

HyperFast performance improvements for write I/O provide for an additional bonus - faster boot ups. To demonstrate this, a natural operating environment was produced, and the Microsoft tools, Xperf and BootVIS, were used to measure boot time.

Testing was done on a brand new system with a 16GB SSD. After installing Windows XP Home edition with service pack 2 from CD, service pack 3 and all updates (to Feb/2009) were downloaded and installed. Microsoft Outlook and Adobe Acrobat were then also installed. The test results are provided below:

Boot Up Performance – xperf.exe tests			
5-run Average	**Out of Box**	**With HyperFast**	**Improvement**
bootDoneViaPostBoot[1]	23923.2	22670.2	5.23%

Figure 5: Benchmarking the time to boot to Explorer and Full Boot (with Microsoft's Xperf)

Boot Up Performance – bootvis.exe tests			
6-run Average	**Out of Box**	**With HyperFast**	**Improvement**
Logon + Service[2]	13.57	12.99	4.3 %
Boot Done[3]	11.06	10.63	3.89%

Figure 6: Benchmarking the time to Logon and start services (with Microsoft's BootVis)

The data shows HyperFast technology improves even out-of-box boot times. Given the above data about increasing free space fragmentation, a strong likelihood exists that boot up times will lengthen proportionate to usage over time. HyperFast will provide faster than new boot up and maintain this speed increase.

[1] This measures to the point where the system is reasonably idle and responsive to user input, the 'PostBoot' phase can be considered complete. Details on Xperf.exe can be found at: http://msdn.microsoft.com/en-us/library/cc305221.aspx

[2] This measures the time required to start Winlogon and any services and applications such as a firewall or antivirus software.

[3] This measures the total time that is taken for XP to Bootup (more info on BootVis: http://articles.techrepublic.com.com/5100-10878_11-5034622.html)

Lifespan:

NAND flash drives have limited erase-write cycles. This fact mandates that any optimization program that may seek to increase performance cannot acceptably do so at the expense of wearing out the drive faster. A tool from BinarySense Inc called SSDLife™ was used to measure the average erase-write activity on the SSD device from various actions. It uses proprietary measures to "score" the erase-write activity from

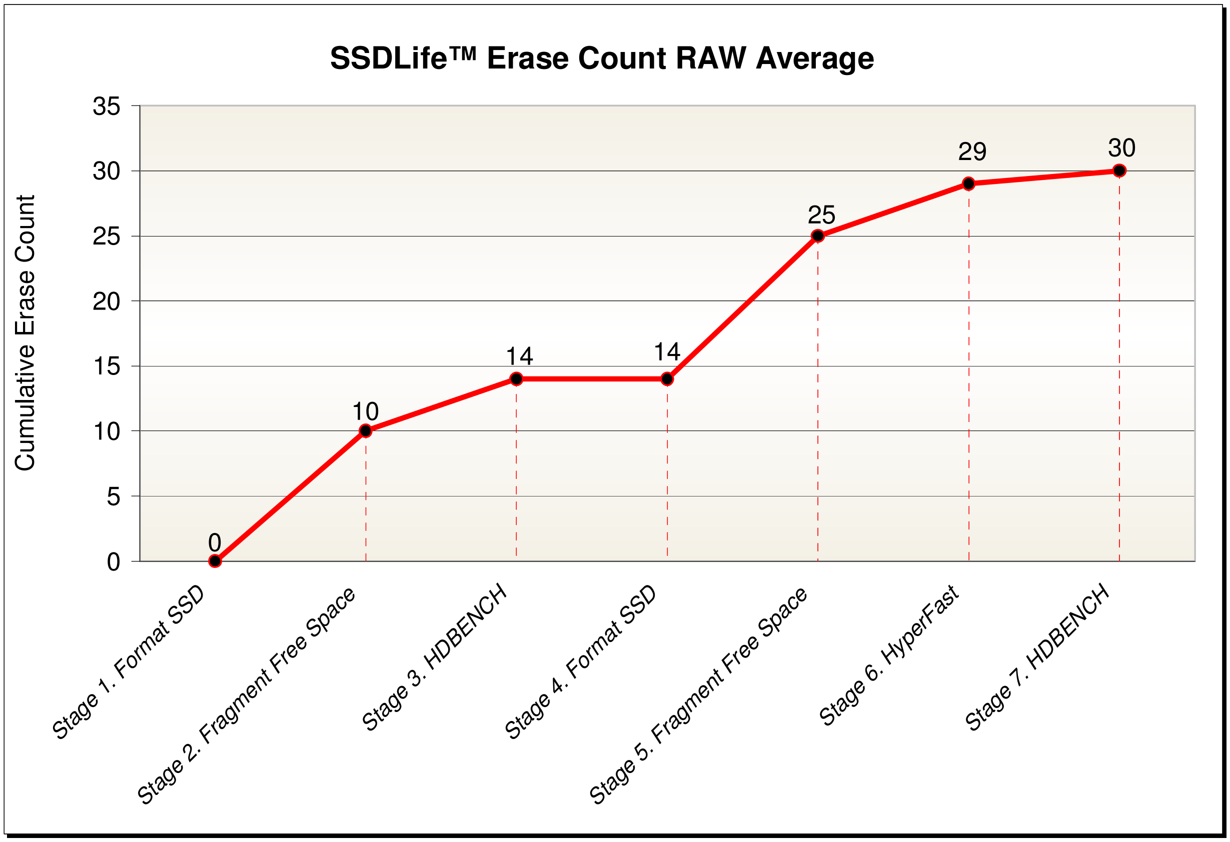

Figure 7: Cumulative Erase activity on 8GB PATA SSD as measured by SSDlife software.

Starting from a freshly Windows' formatted drive (stage 1), fragmenting the free space (stage 2) generates a score of 10 from SSDLife's proprietary scoring system. Then running the HDBench program (stage 3), which creates and reads files to measure hardware performance, generated a count of 4. A reformat (stage 4) incurred no change. A second run of the free space fragmentation routine (stage 5) averaged an Erase count of 9. HyperFast was then run (stage 6) to optimize the SSD. It created Erase activity at an average count of 4 (identical to HDBench in stage 3). HDBench was then run on the optimized SSD (stage 7), and this time only incurred an average Erase count of 1.

This test indicates that, while the SSD optimization techniques employed by HyperFast do incur the occasional increase in erase-write activity, the result of its operation reduces the erase write activity of typical day to day use of the SSD. The net result of using HyperFast will be less erase-write activity on the Flash drive, thereby increasing its longevity.

It should also be noted that in the test case above, the effort required from HyperFast (in stage 6) was the reparation of many months worth of accumulated free space fragmentation. HyperFast specialized algorithms intelligently optimize the drive's performance and maximize its useful life.

Summary:

With typical system use, the out of the box performance of computer systems shipping with solid state storage devices will degrade. This is not due to quality or design of the hardware, but rather the fact that inherent attributes of file systems will manifest over time.

Fragmentation, in this case primarily of free space, is a natural side effect that will occur and accumulate over time with all file systems. As free space becomes increasingly fragmented, performance will suffer proportionately. This leads to increased, excess write activity on the NAND Flash device, diminishing the device's longevity and reducing its performance dramatically, especially with write speed.

In summary, this breakthrough SSD optimization technology from Diskeeper Corporation delivers extended life and greater performance and faster system boot up for the full life of NAND Flash storage devices.

Diskeeper Corporation China Limited
21st Floor, ICBC Tower, Citibank Plaza
No. 3 Garden Road, Central, Hong Kong
www.diskeeper.com/china