SEAGATE

Enhanced Caching Advantage—TurboBoost™ and Advanced Write Caching

Technology Paper

How Seagate® Enterprise drives are once again optimizing performance

Today's enterprise IT professionals evaluate storage drives according to several criteria: warranty, capacity, RAID rebuild features, and so on. But when it comes to performance, nothing has held a candle to solid state drives (SSD) with regard to IOPS and latency metrics. And yet, systems still need the superior capacity and storage density of hard drives (HDD). Storage arrays with both a number of SSDs and HDDs installed can provide an all-inclusive solution, but they entail higher complexity, more management and potentially elevated costs. In the ongoing challenge to strike the best balance between capacity, performance and per-terabyte value, the market has yet to settle on a workable solution.

In seeking out application storage options, IT groups often decide on a desired capacity, then shop for the lowest latency it can afford at that size. Product specifics within any drive solution figure heavily into this decision. Too often, drives struggle to average a response time of under 40ms, and array overhead and increasing queue depths only make matters worse.

When HDD and SSD Join Forces

Seagate opened the door to improvement as engineers and customers began to see the benefits of building NAND flash into HDDs in an assistive role. Widely known as solid state hybrid drives, or SSHDs, Seagate also integrates the technology as TurboBoost in the enterprise market. Rather than construct storage solutions with a mixture of discrete HDDs and SSDs, TurboBoost can often provide the best of both worlds. At the very least, TurboBoost can help to improve general work productivity through greatly improved availability of hot data.

Furthermore, Seagate is about to significantly expand the frontier for high-performance enterprise HDDs with a range of enhanced caching features designed to address three critical objectives, as follows:

1. Avoid as many seeks as possible.
2. Read and write data from the fastest access point (cache).
3. Allow the drive to access and write data in the most efficient way (rotational positioning).

Seagate's methods for meeting these three objectives are now yielding enterprise HDD upgrades that will double or even triple the IOPS performance of conventional storage options without resorting to dedicated SSD caching or expensive full SSD arrays. In fact, Seagate is applying different NAND and DRAM implementation strategies to deliver even higher performance in various enterprise applications, including database, mail and web serving, video distribution, cloud services, and much more.

Workload and Performance

Traditional HDDs predictably respond to higher workloads with slower response times. At 200 IOPS, a drive array might provide a 5ms response, but at 600 IOPS, latency could rocket to 40ms or higher. Fortunately, much of this strain can be alleviated through the benefits of enhanced caching (Figure1).

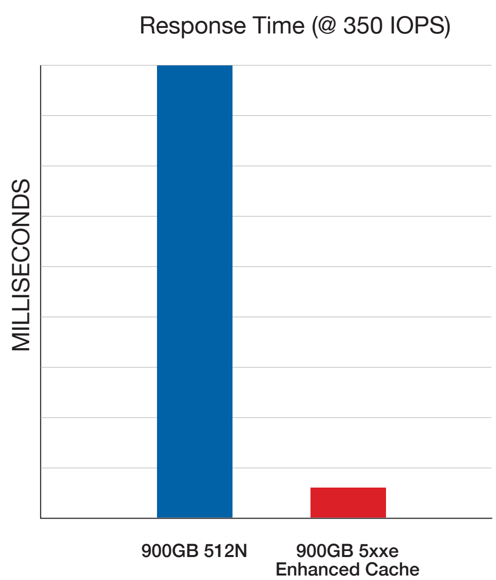

Figure 1. Advanced write caching>10× faster response times

Testing under common industry benchmarks, such as SNIA's hot band workloads using Vdbench, shows that a similar array configuration stocked with Seagate's Enterprise Performance HDDs with enhanced caching don't surpass 5ms until beyond 800 IOPS. In short, the greater the workload, the more effective SSHDs are.

Normally, workload parameters will vary depending on the application and specific datasets. The mixture of random vs. sequential data, reads vs. writes, average file size, block size, and so on can significantly impact performance. Online analytical processing (OLAP, also called data mining and business intelligence), enterprise resource planning (ERP), and online transaction processing (OLTP) constitute the largest application categories that relies more heavily on data reads than writes. Most IT spending flows toward read-intensive apps, which show strain as enterprise workloads increase. This is why TurboBoost was Seagate's major performance enhancement to standard Enterprise Performance HDD design—it addresses the most prominent need in the enterprise storage market.

Conversely, data warehousing and video streaming focus predominantly on writes. Due to the nature of how NAND functions, SSDs and flash caching provide less benefit on write processes than on reads, making the technology less appealing from a performance-per-dollar perspective. The two standout applications that stress write performance—video streaming and data warehousing—have been slower to move beyond traditional HDD designs for this reason, although those applications also crave higher throughput as loads scale.

Lower-demand apps, such as content and project management, tend to perform well enough without NAND flash assistance. However, where reads or writes are intensive, NAND flash caching is gaining ground because regular HDDs can't always meet latency and IOPS requirements, improvement.

Going Beyond the Basic HDD

Conventional drive technology is well-known. Externally, a host controller issues read and write commands to the drive. Within the drive sits two components: main storage (the rotating magnetic platters) and a DRAM cache that acts as a buffer. The process of getting read/write commands executed breaks down into several stages. We can illustrate this by examining the six stages of how a conventional HDD performs a data read.

1. Read command. Put simply, the drive controller receives a read command from the host and begins to execute it.

2. Queue. Ideally, a command can execute as soon as it is received, but in busy environments the odds are that the drive isalready busy performing another task. The new command has to wait its turn in a queue.

3. Seek. When ready to read data from the magnetic media, the drive must position the appropriate read head over the specific tracks containing the desired data. This is known as seek time.

4. Rotational latency. Once the head is over the correct track, the head still needs to wait for the sector within that track to passunder its magnetic sensor. Rotational latency is the average time it takes for a sector to be readable by a head. The faster the platter rotation, the lower the average latency, which is why enterprise-grade, performance-oriented drives feature 10K and 15K spin speeds.

5. Media transfer. As the name implies, this is when the drive reads data from the disk media back to the host.

6. Status. When all of the requested data has been copied from media back to the host, the drive sends status to confirm completion.

Within these six steps, the drive's DRAM plays a critical role. Common uses include command queuing and running write-cache-enable (WCE) to accelerate writes by allowing the drive to store commands before they have been written to long-term media. As soon as data finishes copying into DRAM, the drive signals back to the host that it's ready for more commands.

DRAM excels at fast processing, which is why it serves as primary system memory (RAM). Its drawback, of course, is that it is volatile memory. If power gets cut during a write operation, then the host acknowledges that data has been successfully sent (because the hard drive signaled that it received it in DRAM), even though the data has yet to be written to magnetic media. With no warning or alert, this in-transit data simply vanishes.

Conversely, NAND flash memory is non-volatile albeit considerably slower than DRAM. This makes it less suited to the frequent buffering tasks required of hard drives. But consumers are willing to accept the slightly lower performance of NAND as long as they do not have to worry about increased data loss risk. On the other hand, enterprises tasked with safeguarding client data often need a more balanced approach.

Taking the TurboBoost Step

As discussed earlier, TurboBoost involves the addition of NAND flash memory into a conventional hard drive. Specifically, drives such as Seagate's Enterprise Performance models with TurboBoost integrate of conventional multi-segmented cache as well as a small amount of eMLC NAND, which features significantly more robust endurance than consumer-grade MLC. Seagate testing has determined that a modest amount of NAND hits a value sweet spot in the majority of current target applications, delivering the most performance benefit at a cost low enough for Seagate to integrate into the overall cost of the drive.

For nearly a decade, engineers have worked to refine how and when data gets copied into an SSHD's NAND. In fact, the algorithms governing this process make up some of the secret sauce that clearly differentiates one SSHD brand and drive family from another. Ultimately, though, the underlying idea remains constant: The data most frequently sought by the host, so-called hot data, gets copied from magnetic media into NAND. Subsequent requests for this data by the host can then be fulfilled far faster from flash memory than from rotating magnetic platters. As the cache fills up, the least *hot* files will retire from NAND to make room for more data while the original files remain accessible from the disk media if needed.

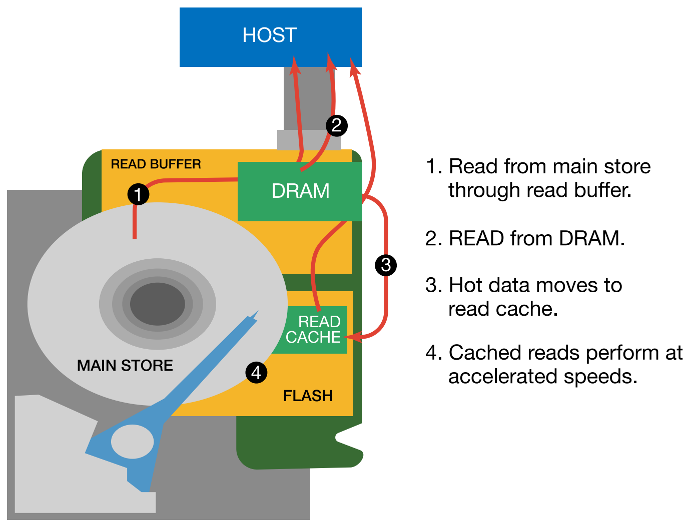

Figure 2. How TurboBoost works

With TurboBoost technology, the seek and rotational latency parts of a conventional drive's read process often vanish, because if data resides in NAND, there is no spinning media to navigate. Read command and status processing remain about the same in duration, but queuing time is shorter, as is transfer duration. The real world benefits of TurboBoost will obviously depend upon application and workload specifics, but in general, keeping copies of the most frequently-used data in flash memory will accelerate read operations by several times. In particular, applications that rely on frequent small block transfers stand to benefit the most, since these will suffer most heavily from seek and latency delays.

OLAP, ERP and OLTP may be the best applications for TurboBoost enhanced caching, but they're not the target segments showing growth. To that list, we should add object-based storage and virtualized server and desktop platforms. As StorageReview explains in a recent independent review, the Enterprise Performance TurboBoost drive easily outperforms, both in throughput and latency, and across a wide range of enterprise applications. StorageReview notes that TurboBoost provides "a very noticeable performance boost over standard HDDs and the SAS interface makes it more capable than the low-cost SATA SSDs that compete on dollar/GB."

That said, results from StorageReview and others confirm that the margin of improvement for read tasks with TurboBoost far outstrips writing. To address the specific market needs of write-focused applications, other optimizations must come into play.

Optimize Further With Advanced Write Caching

Although TurboBoost is a relative newcomer on the scene, DRAM buffers have been a mainstay in hard drive design for many years. The market, however, contains even more niches in need of storage optimization, so Seagate recently released further enhancements for advanced write caching among its enterprise drive portfolio. Expect enhanced write caching to pay maximum dividends in applications such as data warehousing, digital content creation and centralized video surveillance.

Mirrored DRAM (mDRAM)

mDRAM is a segment of the drive's DRAM that has a copy of data written to media cache. This is so that the drive can actually read the data back to the host or write the data to the main disk store if the data is no longer consider hot. This saves on rotational and seek delays inherent in navigating to media cache positions—effectively preventing a double seek of both media cache and the main store. Moreover, the mDRAM sorts write commands to take full advantage of rotational positioning at any given time.

NVC-Protected Write Cache

Back electromotive force (back EMF) is a voltage that appears in opposition to the current flow generated by a motor's coils moving through a magnetic field. A hard drive can use the back EMF of its spinning down motor immediately after a power failure to write the data contained in DRAM to a block of NVC flash memory. In essence, the drive uses its dying breath during power failure to pass on DRAM data to non-volatile NAND. With NVC-protected write cache, a drive can retain the performance benefits of operating in write cache disabled (WCD) while still having write caches be protected (Figure 3).

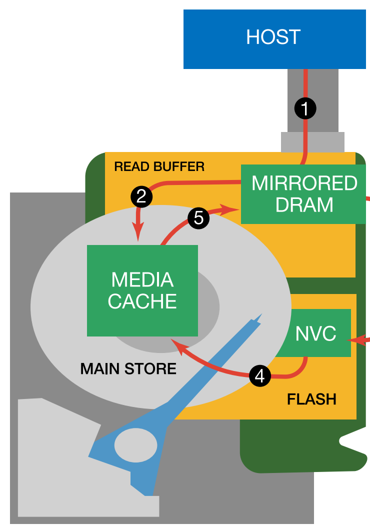

1. Data writes into NVC-protected DRAM; then to disk.
2. Writes coalesce in cache and sequentially write to media cache a copy remains in DRAM, (large transfers go directly to Main Store).
3. If power is lost, DRAM writes to NVC for data retention.
4. On power-up, NVC data writes to media cache.
5. DRAM restores from media cache.

Figure 3. How advanced write caching works

Enhanced Caching—Incorporating Technology With Immediate Read and Write Benefits

Seagate's enhanced caching feature consists of advanced write caching to improve write performance, and TurboBoost for improved read performance. Advanced write caching is now featured on many Seagate nearline hard drives, while TurboBoost read caching is standard on Seagate's mission critical HDD models. By integrating the advanced write caching on the Enterprise Performance 15K HDD, Seagate can offer up to a 100% improvement in random write performance over the previous generation—
(RW/RR : 720/400 IOPS).

Seagate plans to unleash a range of enterprise drives that integrate the enhanced caching feature in varying combinations.

Businesses and enterprises will be able to choose the exact drive best able to fit their workloads. Furthermore, Seagate will provide communication and tools for creating that optimal pairing.

We can already see early evidence of the competitive advantages of Seagate's enhanced caching feature. For example, Seagate's Enterprise Performance 15K 2.5-inch HDD outperformed HGST's 15K 600GB by approximately 3× across mixed workloads. This is due in large part to Seagate's implementation of the enhanced caching feature (Figure 4).

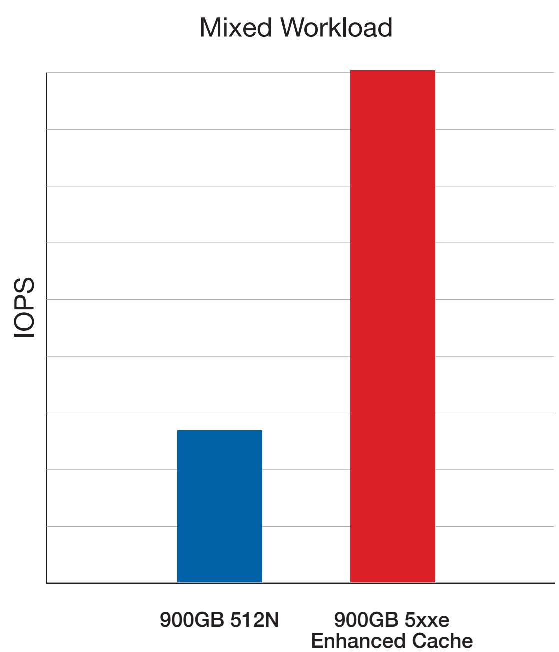

Figure 4. TurboBoost with integrated NAND flash performs up to 3× faster

The more your organization relies on melding storage capacity with performance, the more you will want to investigate enhanced caching and examine which specific features match your workload demands. Assess the value benefits compared to a more complex HDD/SSD array solution. Determine ROI estimates for enhanced cache across your current and coming deployments, and learn more about taking next steps at < Enterprise Performance 15K HDD.>

seagate.com

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16–18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00