

## MAINSTREAMING HIGH IO PERFORMANCE WITH FLASH CACHE

DECEMBER 2012



Servers are growing more and more powerful but decades-old storage controller technology has not kept pace. This puts storage at an extreme disadvantage just as companies are growing huge infrastructures of virtual and physical systems, applications of all types, and very large volumes of primary data.

Advances in solid-state drives have helped by introducing higher performance/lower latency into the computing stack, and customers do not necessarily have to sacrifice simplicity or budgets for performance improvements. However, these SSD products – which may be at the server, network or storage layers -- can be very expensive and complex to deploy and manage even at the enterprise level, let alone the mid-market. Customers want and need the performance boost from SSD products but they are asking if the improvement is worth the expense, management complexity, and confusion around different SSD offerings. The answer is: it depends.

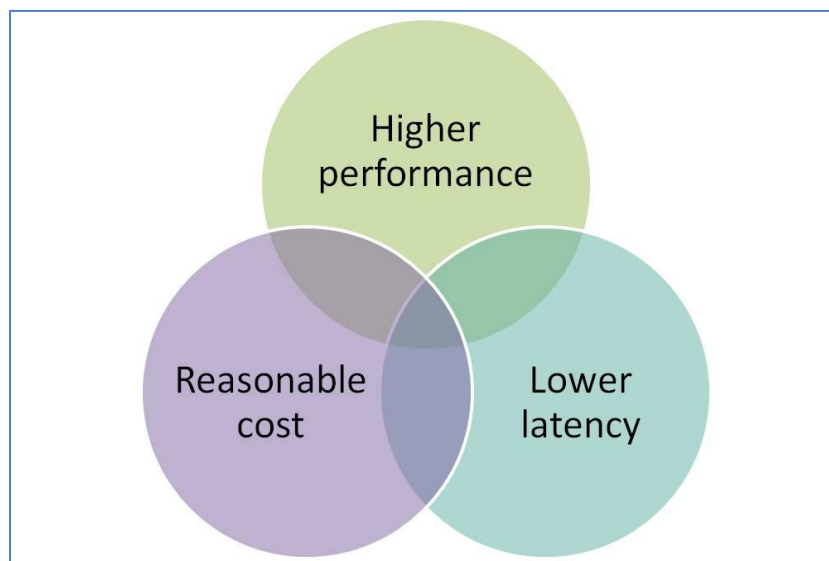
### *SSD Adoption Factors*

Factors	Notes	Our Take
Where is the SSD located?	SSD may be placed up and down the computing stack at the server, networking or storage levels.	SSD can be useful at the server level for single Tier 1 applications, but the storage cannot be shared. Most applications use multi-tenant storage arrays so it makes more sense to locate flash at the shared storage level.
What is the storage array's architecture?	All-flash arrays and hybrid arrays using flash cache or flash tier are different ways of providing increased performance.	All-flash arrays are expensive and have relatively small capacity. Hybrid flash arrays combine flash and disk.
Does the hybrid array use flash cache or tiering?	Flash caching or tiering optimize read/write performance, which cuts down on the purchase of additional flash or high performance spinning disk.	We like flash caching over tiering because of cache's management simplicity and data redundancy. We particularly like flash cache combined with DRAM for large incoming data and worksets that fit into cache.

### Why Flash Matters to Applications

Storage-based SSD solves serious performance issues by providing fast IO processing at the storage controller, making it ideal for high performance and low latency storage requirements. The wrench in the works is that all-flash arrays are quite expensive and small capacity compared to hard disk, making them out of reach for most mid-sized businesses, and unnecessary for the majority of data center applications. Some critical Tier 1 application environments can take advantage of dedicated all-flash arrays or even card SSDs at the host to dramatically accelerate applications whose service levels are measured in microseconds.

Most mainstream applications do not need the stratospheric IOPS numbers advertised for all-flash storage systems, but there are still crucial performance requirements for mission-critical applications such as databases and OLTP applications, SharePoint, Exchange, virtual servers, and virtual desktops. While businesses of all sizes must maintain high random IO performance and low latency for these applications, they cannot justify paying the premium prices for dedicated all-flash storage.



There is a way to apply the high SSD performance with reasonable cost per GB: hybrid storage systems that combine SSDs and spinning disk to provide high performance storage at a reasonable cost. In this environment, the flash-based storage controller provides intelligent storage management, SSDs provide fast performance for more active data, and back-end disk houses the majority of data. Yet even in this category, systems differ from one

another in several important respects. One of the most important differentiations is in fact one of the least obvious: does the array use cache or tiering to deliver active data into the SSD layer? At first glance these two approaches do not look too different, but in fact their differences over time can be substantial.

- **Automated Tiering.** Data tiering has been around for years and is the process of matching data value to storage costs. Traditional tiering moves disk-based data down from primary to nearline, archival and tape tiers for economical long-term retention. Primary tiering using SSD has a very different meaning. SSD tiering moves data back and forth between Tier 0 SSDs and Tier 1 fast spinning disk. Movement occurs in response to application usage, a manual command, or a schedule where data sets are periodically moved between the SSD layer and spinning disk. The data moves, not copies, so must be highly protected using some form of SSD RAID. Additional RAID calculations may impact performance and add to the cost with additional

SSD capacity. These are not granular moves; GB-size blocks and file volumes will move even if the actual data request is very small.

- **Caching.** Caching has also been around for a long time but SSD caching changes the equation considerably. Rather than tiering, which requires administrators to understand IO usage and set policies, caching uses predictive algorithms to copy data to the SSD layer for fast IO processing. The storage array may do write caching as well as read caching; we much prefer to see both. Write caching buffers incoming writes and efficiently aggregates them before doing disk writes. Using SSDs for the read cache enables the array to have a much larger cache, which increases the amount of data the cache can store and accelerate.

### **Advantages of Caching**

Flash caching drives dramatic performance gains by keeping active data copies on very fast flash cache. This saves money since customers do not have to increase spindles to increase performance. Some customers can even replace expensive SAS with SATA because they gain so much performance from the SSD cache. Customers also save money because flash caching can be effective with a relatively small amount of flash and does not require heavy management overhead; they can leave optimization to the caching algorithms.

Caching's predictive learning algorithms takes that problem out of administrators' hands by learning and applying the optimal placement for best application performance. Larger SSD capacity, granularity, simplicity and redundancy also work into the caching equation:

IO placement uncertainty is not at all limited to the mid-sized business. Few administrators even at the enterprise level really know how to manage data placement by IO performance, even when their critical applications are suffering from related performance problems. The best solution? Let the intelligent cache do what it was designed to do: optimize IO performance and data placement.

- **Larger flash capacity.** Caching has benefited from important advances in flash SSD sizes, which have reached TB capacity, entire large worksets can fit into cache for very fast IO processing.
- **Granularity.** Cached data copies are highly granular. Flash tiering moves large chunks of data even if a just small block or byte has changed. But caching can copy large or highly granular data. This minimal data copy dramatically increases performance over large data movements between a flash tier and backend disk.
- **Simplicity.** Management simplicity saves on ongoing expenses, which is key to cost savings over the life of the array. Letting the storage system provide good caching algorithms is preferable to requiring IT to figure it out data placement and matching best performance for an application's needs.
- **Fast reads and writes.** Read and write flash caching greatly accelerates performance over spinning disk drives. Reads are considerably faster than writes, but caching vendors will usually employ write acceleration technologies just as writing from DRAM, aggregating writes, and employing logs.
- **Accelerates overall storage performance.** All applications sharing the SSD storage will increase performance and decrease latency. Benefits spread to all end-users.

- **Real-time vs. periodic.** Most automated tiering solutions do not operate in real-time but offer manual commands, factory-set data movement schedules, or policy-controlled data movement between tiers. Flash cache operates in real-time to automate and optimize copy movement between tiers as the workload changes.

Let's look at an example of a hybrid storage array with SSD-based caching, the Nexsan NST5000. The array combines gigabytes of DRAM, terabytes of high capacity flash drives and SAS or SATA for long term spinning disk capacity. Nexsan calls its cache technology FastTier, which performs reads and writes in its high-speed DRAM and flash layer for extremely high performance and low latency. Frequently used data can granularly stage to the cache for fast processing. The NST5000 does not need to RAID its SSD read cache because the data always lives on the underlying RAID'ed spinning disks. Performance for sequential files is stellar; performance for random IO files – by far the harder task – hits 100,000 CIFS op/sec at 1.23ms latency.

SSD cache helps to solve the big business problems of poor performance and high latency for mainline applications, without a large per/GB price tag. When utilized within hybrid storage arrays, SSD-based cache consolidates high performance storage onto a shared storage system that costs a fraction of all-flash and in many cases just slightly more than an all disk based system, and greatly multiplies IO performance over traditional storage arrays.

---

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.