NEXSAN

HYBRID STORAGE WITH FASTier® ACCELERATION TECHNOLOGY

Nexsan's FASTier acceleration technology leverages the power of solid-state to accelerate the performance of the underlying spinning disks by a factor of up to 10X. The net result is a radical new cost-per-I/O, allowing organizations to do more than ever before in virtual and cloud environments.

TECHNICAL BRIEF

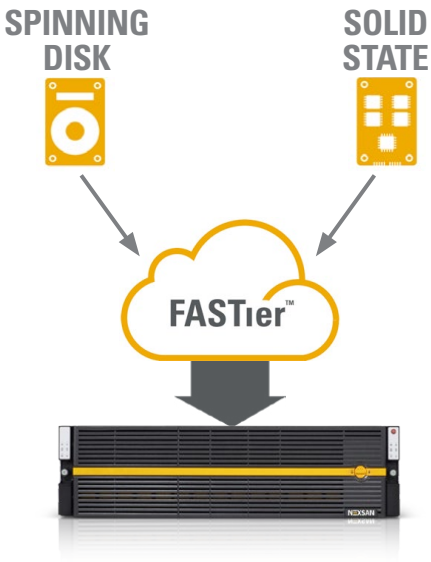

THE SPEED OF SOLID-STATE AT THE PRICE OF SPINNING DISK

INTRODUCTION

The onset of virtualized environments and cloud-based computing has caused traditional disk-based storage to struggle in keeping pace with the new I/O pressures being placed on it. As the number of virtual machines grow, even the fastest traditional storage with the most powerful controllers are simply not fast enough to scale with the seemingly endless performance demands of a sprawling environment of virtual servers and desktops, mission-critical databases and email servers. The result has been a storage performance bottleneck that keeps companies from fully leveraging the true application power and cost-effectiveness of their virtual and cloud infrastructures.

Enter Solid-state

Solid-state performance alleviates storage performance choke-points... but at a price and capacity that is unreasonable for the average organization. In the rush to exploit an emerging marketing opportunity made possible by SSDs, some vendors are providing SSD-only arrays. However, the exorbitant cost and capacity constraints of an SSD-only array have relegated its use to extreme situations. Some vendors have introduced SSDs into their systems with automated tiering software in which data can be migrated from SSD to SAS and/or SATA drives. However, this operation can significantly reduce the performance of the system as data is moved from one tier to another. Auto-tiering presents additional problems as organizations must purchase an inordinate amount of SSD to store their most frequently accessed data because they are housing data rather than caching data. Once data is moved out of the SSDs, it is subject to the slower speed of the spinning disks.

Enter Hybrid Storage

Nexsan's FASTier acceleration technology has innovatively combined the performance of solid-state with the capacity and cost effectiveness of traditional disk storage. Imagine a Hybrid Storage system as an orchestra. If rotating media is the percussion section and solid-state is the brass section, FASTier is the conductor that makes them work together in harmony to provide the music - a Hybrid Storage system that delivers the speed of solid-state along with the capacity and cost benefits of spinning disk. FASTier's advanced software architecture and algorithms manage the data that is being read and written, leveraging a modest amount of solid-state to greatly accelerate the performance of the underlying spinning disks by a factor of up to 10X or more.

In this scenario, solid-state is judiciously used to enhance traditional disk-based storage via FASTier to provide new levels of performance where cost and capacity are primary concerns. As a result, a Hybrid Storage system (like Nexsan NST5000)

CACHE LOCATION OPTIONS

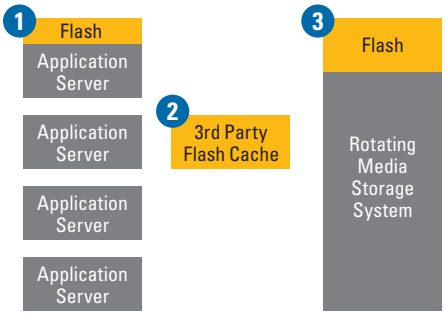

easily eliminates the performance bottlenecks in the most demanding virtual and cloud environments.

This technical brief will survey the landscape of flash-based caching approaches, identifying their strengths and weaknesses, and explore Nexsan FASTier in the Nexsan NST5000 Hybrid Storage system.

CACHE SIZE

Caching techniques are not new to storage systems or application servers. Reads and writes are typically held in high-speed DRAM providing needed acceleration for the underlying hard drives. Frequently requested blocks can be read from the spinning disks and held for a period of time. Smaller writes are also held in DRAM where they are aggregated into larger writes and RAID checksum calculations can be performed before they are pushed to the spinning disks.

FASTier acceleration technology in the NST5000 storage system utilizes up to 192GB of DRAM and up to 2.8TB of 100GB or 200GB flash-based SLC SSDs. The NST5000 is a Hybrid Storage system that contains up to a petabyte of spinning disk storage. FASTier is ideal with that kind of capacity as it accelerates overall storage performance by up to 10X. If the working sets of the application servers fit entirely into the SSD cache, performance can climb even more dramatically.

CACHE LOCATION:
Application Server, Storage System, or In-between

There are three places to introduce cache to accelerate storage performance: (i) in the application server, (ii) in the storage system, and (iii) in an appliance introduced between the application servers and the storage system. Each has its advantages and challenges.

When cache is introduced in the application server itself, it lives on the high speed PCIe bus for maximum bandwidth. Still, since it is a cache, the speeds at which it can be filled and drained must be balanced, lest it remains unfilled, or becomes completely filled. So, bandwidth to the storage system has to be considered. Multiple 8Gb Fibre Channel, 10Gb iSCSI, 24Gb SASx4, or 40Gb Infiniband interconnects may still be needed.

There are four disadvantages to application server-based caching:

- Cache is typically not resilient to a single point-of-failure. It can only be used to cache reads. Writes must be pushed all the way to the storage system before they can be acknowledged safely to the application. The inability to cache writes has a pronounced negative impact on VDI performance and other applications.

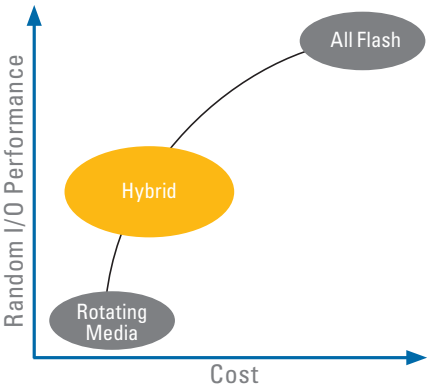

- Custom proprietary device drivers need to be installed onto the application servers, a practice that is generally unacceptable or highly questionable to IT administrators.
- Application server-based caches are, by design, a dedicated resource to a single server and thereby reduce operational flexibility and the ROI on the purchase.
- Perhaps most importantly, server CPU cycles are used to run caching algorithms, detracting from the total power available to run applications.

When storage systems cannot be outfitted with a sufficiently sized cache, a third system can be introduced in between the application servers and the storage system. This is the least beneficial architecture of all, and usually the most expensive.

There are four disadvantages to interstitial appliance-based caching:

- Management is difficult since the appliance cache must be managed separately from the application servers and the spinning disk storage system.
- Should it fail, all the writes that it was caching may be lost. As a write-through cache, only reads are cached, so writes must go all the way through to the spinning disk storage system before they are acknowledged. The result is a significant overall performance loss.
- Support contracts are complicated as storage vendors may simply blame the interstitial appliance in any troubleshooting scenario.
- Many IT administrators avoid interstitial appliance-based caching because the risks are too high and reliability is questionable. The best place for cache to get the greatest storage performance advantage is in the storage system itself.

There are five advantages to storage system caching with FASTier:

- FASTier storage-based caching offers a fault-tolerant architecture, enabling both reads and writes to be cached, which is especially important for supporting virtual desktops.
- No proprietary or unproven device drivers need to be installed on the application servers.
- Since the cache accelerates overall NAS/SAN performance, all application servers that use the storage system will benefit from it.
- The caching algorithms utilize the CPU power and DRAM in the storage system, so application servers are not bogged down.
- Management of the storage system-based cache is usually integrated with overall storage system administration and with other important storage system features such as thin provisioning, online capacity expansion, snapshots and replication.

NST HYBRID STORAGE ARCHITECTURE

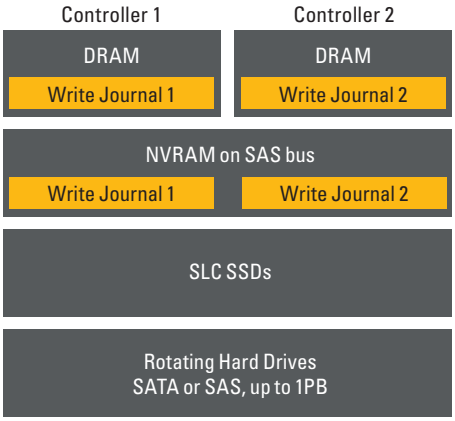

Flash-only Storage Systems

Rather than use flash-based devices as a cache for spinning disks, some storage systems are flash-only. While the performance numbers are high for such systems because the entire working set fits into flash, the costs are exorbitant while the capacities are extremely low, typically under 50TBs as compared to Hybrid Storage systems that leverage spinning disks to hold up to 1PB or more. What's more, flash-based SSDs have a much shorter lifespan compared to spinning disks. And when they are configured into a RAID set, the SSDs may experience two writes for every application write – one to store the data, and one to store the RAID5 or RAID6 checksum, further decreasing the lifespan of the SSDs.

Cache Hardware Architecture

Not all hardware architectures are alike for flash-based caches. Some are not resilient against a single point-of-failure. Therefore, they can only be used as a write-through cache where reads can be cached but writes experience long delays as they are committed to the spinning disks before being acknowledged. FASTier is fully fault-tolerant. It can cache both reads and writes for maximum application performance acceleration.

FASTier Overview

FASTier uses DRAM, NVRAM and flash caching to greatly accelerate reads and writes in concert with the underlying spinning disks. DRAM and NVRAM are important parts of the hardware architecture because they are 20X or more the speed of flash and they do not wear out. FASTier uses flash for both read and write caching. In write caching, the transactions are protected in the DRAM/NVRAM-based write journals until they make it all the way to the spinning disk RAID set, so the flash devices do not need to be in a RAID set. Further, the read cache flash does not need to be configured into RAID sets. Should a read cache flash device fail, it is mapped out and the read is re-issued to the underlying spinning disk. NST5000 Hybrid Storage minimizes the wear on flash-based devices by clever use of DRAM, NVRAM to prolong the overall life of the system.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com

"

FASTier realigns cost, capacity and performance by offering the speed of solid-state with the capacity and price of spinning disk."

NEXSAN

Cache Software Architecture

The key to FASTier is the software architecture and algorithms that intelligently and automatically use solid-state devices to accelerate the performance of the underlying drives. The FASTier write journal uses the same technique that has been used in large enterprise-class databases such as Oracle, DB2, or Microsoft SQL Server, where writes to the database come in as transactions, are placed into the write journal immediately, then acknowledged to the writer. Later, they are picked up from the write journal and applied to the underlying hard drives. This same technique is used in journaled files systems such as Microsoft's NTFS and Apple's HFS Plus journaled file system. Another FASTier algorithm uses DRAM to aggregate smaller writes before sending them to the underlying spinning disks, which increases the ability of the spinning disks to absorb the data efficiently. Opportunistic read-aheads of blocks for applications exhibiting proximate read access patterns also increase system performance. And holding transactions in the write journal until they make it all the way to the underlying drives increases cache performance and minimizes flash wear.

Cache Management Matters

FASTier management is performed within the GUI of the NST5000 storage system. The only decisions are the size of the FASTier read and write caches, the number of FASTier caches per system, and the storage pools that each FASTier cache will be associated with. This level of simplicity is in stark contrast to the complexity of caching features on competing storage systems, server-based caches, or third-party appliance-based caches that live between the application servers and the storage system.

Overall Storage System Considerations

Nexsan's FASTier acceleration technology utilizes solid-state devices and spinning disks to create a Hybrid Storage system which breaks through the barriers of traditional spinning disk storage. The resulting performance, cost and capacity benefits are far beyond alternative caching architectures and flash-only storage systems. FASTier changes everything. A FASTier enhanced Hybrid Storage system, such as the NST5000 with 7200RPM SATA or NL-SAS drives, can deliver the performance of a system with 15K SAS drives, at one third the cost and in one third the footprint, while using one third the power and cooling. And because FASTier on the NST5500 can be configured up to 2.8TB, it can hold entire working sets to deliver up to 10X the performance of spinning disks and beyond. Alternatively, an NST5000 Hybrid Storage system outfitted with SAS drives and FASTier delivers industry-leading performance and price-per-IOP to power the most demanding virtual computing environments and databases.

NEXSAN 1445 Lawrence Drive, Thousand Oaks, CA 91320 | p. 866.4.NEXSAN f. 805.418.2799 | www.nexsan.com