NEXSAN

HYBRID STORAGE WITH FASTier™ ACCELERATION TECHNOLOGY

Nexsan's FASTier™ acceleration technology uses advanced software architecture and algorithms to leverage the power of solid-state to accelerate the performance of the underlying rotating media by a factor of 3X-10X or more. FASTier is the core performance technology that makes Nexsan NST5000™ unified hybrid storage systems ideal for virtual and Cloud environments as well as traditional environments.

TECHNICAL BRIEF

NEXSAN

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. 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 easily 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 market opportunity made possible by SSDs, start-ups are providing SSD-only storage solutions, but their use has been relegated to the very fringe where performance outweighs capacity and price considerations. Traditional vendors introduced SSDs into their systems and with automatic tiering software data can be migrated from SATA or SAS drives to SSDs. However, this operation reduces the performance of the system to such degree that it is best run at night. Auto-tiering presents additional problems. Organizations must purchase an inordinate amount of SSD to store all their highest frequency data, are not able to leverage the fastest solid-state devices for capacity reasons, and once data is moved out of the SSDs it is subject to the speed of the rotating disks.

Enter hybrid storage

Nexsan 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 SATA 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 SATA/SAS. The advanced software architecture and algorithms of FASTier manage the data that is being read and written and leverages the power of solid-state to accelerate the performance of the underlying rotating media by a factor of 3X-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 still primary concerns. By properly leveraging the power of solid-state, hybrid storage systems (like the Nexsan NST5000) easily eliminate the performance bottlenecks in the most demanding virtual and cloud environments.

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

This technical brief will survey the landscape of flash-based caching approaches, identifying their strengths and weaknesses, and explore Nexsan FASTier acceleration technology, which is utilized in the Nexsan NST5000 hybrid unified storage systems, in more detail.

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 rotating 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 rotating drives.

The size of the DRAM cache for a midmarket storage system has traditionally been in the range of 4GB to 16GB. Once the cache is full, the speed of the system degrades to the speed of the underlying rotating drives. There are two solutions to this problem – use more DRAM and/or other flash-based devices. FASTier acceleration technology in Nexsan NST5000 storage systems utilizes up to 192GB of DRAM which is expandable up to 2.8GB of 100GB or 200GB flash-based SLC SSDs. The NST5000 is a hybrid storage system that contains up to a petabyte of rotating disk storage. FASTier is ideal with that kind of capacity as it accelerates overall performance of the storage system 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 are still needed.

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

NEXSAN

FASTier ARCHITECTURE

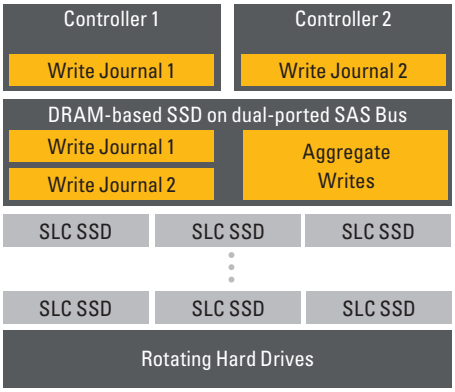

FASTier Software

- Fully Fault-tolerant Architecture
  - › Both reads and writes can be cached
- Write Journal
  - › Proven technique used by the largest databases in the world
  - › Logs writes, then later writes are applied to the rotating media
- Opportunistic read-ahead into flash-based SLC SSDs
- Aggregate smaller writes into bigger ones in the fault-tolerant DRAM to optimize writes to underlying rotating media
- Easy to set up, without the need for any manual intervention
- Per storage pool to focus the power boost on a subset of disks
- Can have more than one FASTier per system

Solid-state Devices in the NST5000

- DRAM-based SSDs on dual-ported SAS bus
  - › Push to onboard flash on power-down
  - › Never wear out, and are faster than flash
  - › Can be accessed by both controllers
- 100/200GB SLC Flash-based SSDs
- From 100GB to 2.8TB to hold entire working sets

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.
- Custom proprietary device drivers need to be installed onto the application servers, 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.
- 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 five disadvantages to interstitial appliance-based caching:

- Management is difficult since the appliance cache must be managed separately from the application servers or the storage system.
- Should it fail, all the writes that it was caching may be lost making it a weak link.
- When set up as a write-through cache, where only reads are cached and writes go all the way to the legacy storage system before they are acknowledged this results in a significant loss of overall performance.
- The storage system that could not take SSD cache is clearly too old, so band-aid fixes to its aged architecture are dubious at best.
- Complicates support contracts 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 four advantages to storage system caching with FASTier:

- FASTier storage-based caching offers a fault-tolerant architecture, enabling both reads and writes to be cached.

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 SATA/SAS.”

- No proprietary or unproven device drivers need to be installed on the application servers.
- The cache is SAN-friendly because 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

Flash-only Storage Systems

Rather than use flash-based devices as a cache for rotating 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 systems that leverage rotating drives to hold up to 1PB or more. What’s more, flash-based SSDs have a much shorter lifespan compared to rotating drives. 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. FASTier uses redundant, dual-ported DRAM module for write caching in most instances, which is 20X or more faster than flash while not wearing out. When flash is used for write caching, the write transactions are protected in the write journal until they make it all the way to the rotating disks in a RAID, so the flash devices do not need to be in a RAID. Further, the read cache SSDs do not need to be configured into RAID sets either. Should a read cache SSD fail it is mapped out and the read is re-issued to the underlying rotating drive. Nexsan NST5000 hybrid storage systems minimize the wear on flash-based devices by clever use of DRAM and rotating disks to prolong the overall life of the system.

Cache Hardware Architecture

Not all hardware architectures are alike for flash-based caches. Some are not resilient to 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 rotating disks before being acknowledged. FAStier is fully fault-tolerant. So, it can cache both reads and writes for maximum application performance acceleration. FASTier also uses two kinds of solid-state devices: DRAM and flash.

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

NST5000

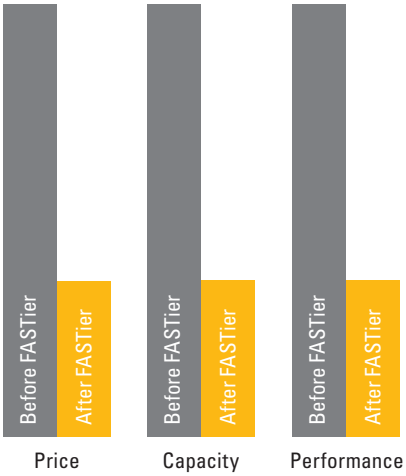

Price
FASTier changes the traditional I/O ratio by accelerating underlying SATA/SAS by up to 10X.

Capacity
By accelerating SATA speed, organizations get new levels of capacity at SAS performance.

Performance
Performance leverages solid-state to accelerate the underlying SATA/SAS by up to 10X.

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 database tables that reside on 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 rotating drives, which increases the ability of the rotating drives to absorb the data efficiently. Opportunistic read-aheads of blocks for applications exhibiting sequential read access patterns also increase system performance. And holding transaction 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 Nexsan 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 FASTier acceleration technology utilizes solid-state devices and rotating drives to create a hybrid storage system which breaks through the barriers of traditional rotating drive 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 Nexsan 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 in size, it can hold entire working sets to deliver unprecedented performance from three times the performance of the rotating drives and beyond.

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