

TP-543 • From: Global Product Marketing • September 2005

Technology Paper

Optimizing Storage With SAS: Beyond the 10K Compromise

Introduction

Enterprise storage is entering a new era of cost-effectiveness and efficiency, driven by the growing movement towards specialized storage solutions. The value proposition of this application-focused storage is straightforward: Optimize price/performance by matching the storage device to the specific characteristics (quantity, needed availability, etc.) of the data. Naturally, such an approach requires multiple device types, each cost-effectively performing its respective storage duties.

Serial Attached SCSI (SAS) is a fundamental enabling technology in this specialization revolution, seamlessly integrating multiple types of storage devices in one enterprise-class infrastructure. SAS compatibility with Serial ATA (SATA) makes it possible to deploy both online (performance) disc drives and nearline (capacity) disc drives on a common SAS backplane in a single enclosure. Coupling this cost-effective infrastructure with optimal online and nearline drives ensures maximum storage value.

Over the years, 3.5-inch 10K-RPM disc drives have performed admirably in a broad variety of storage applications. But now these generalized, “jack of all trades” drives are being supplanted by more specialized, application-specific devices. Eclipsed in both raw performance and price/performance by 15K-RPM SAS online drives (and soundly trounced in cost/GB by SATA nearline-ready drives), 10K SAS drives are a costly compromise, out of place in today’s more efficient storage environment.

Bottom Line: 15K SAS Delivers More

While the factors underlying the outstanding efficiency of 15K SAS storage systems are noteworthy (and covered later in this analysis), it is the significant real-world benefits these systems provide that compel immediate discussion.

A common SAS deployment might entail a 2U, 12-bay enclosure/rack space filled with 3.5-inch disc drives. In order to address both online and nearline applications, the backplane would typically contain a mix of SAS drives and SATA drives. As shown below, the superior speed of 15K SAS drives not only delivers faster throughput, it also frees up enclosure space for more nearline drives. The net result is *both higher online performance and greater nearline capacity* (see table 1):

The 15K Advantage: More Performance, More Capacity					
Storage Configuration: 2U w/12 3.5-Inch Disc Drives	Online (SAS) IOPS per Drive	Online (SAS) IOPS, Total	Nearline (SATA) Capacity, Total	Nearline (SATA) Cost/GB (est.*)	15K SAS Advantage
15K SAS/Nearline SATA vs. 10K SAS/Nearline SATA					
5 - 15K SAS Drives (146GB) 7 - 7200 SATA Drives (400GB)	55.1	275.5	2800GB	\$1.00	11% More Online Performance, 17% More Nearline Capacity
6 - 10K SAS Drives (146GB) 6 - 7200 SATA Drives (400GB)	41.4	248.4	2400GB	\$1.00	
15K SAS/Nearline SATA vs. All 10K SAS					
5 - 15K SAS Drives (146GB) 7 - 7200 SATA Drives (400GB)	55.1	275.5	2800GB	\$1.00	11% More Online Performance, 56% More Nearline Capacity at 75% Lower Cost
6 - 10K SAS Drives (146GB) 6 - 10K SAS Drives (300GB)	41.4	248.4	1800GB (SAS)	\$4.00 (SAS)	
<small>Test results from IBM xSeries 336 with EXP 400, general transactional workload, RAID 5. *Source CDW, August 2005.</small>					

Table 1. 15K SAS delivers more performance and capacity

Should more nearline storage be required, a combination of four 15K SAS/8 SATA drives would yield moderately lower IOPS (220.4) than the reference 10K/SATA configuration, but offer **33 percent greater nearline capacity** (3200 Gbyte). Conversely, when online speed is the top priority, a mix of six 15K SAS/6 SATA drives results in **33 percent higher IOPS** (330.6) than the 10K reference system, while providing the same nearline capacity.

For IT managers seeking to address both online and nearline applications with 10K SAS drives, the associated performance and capacity figures are even more eye-opening. Not only does the 15K SAS/Nearline SATA combination yield higher online IOPS than an all-10K SAS solution, it also delivers *56 percent more nearline capacity at 75 percent lower cost*.

Regardless of the 10K SAS configuration under consideration, the net result is that 15K SAS solutions provide significantly lower total cost of ownership (TCO). Delivering greater performance with fewer drives than its 10K counterpart, 15K SAS reduces storage costs such as infrastructure, IT maintenance support, data center footprint, and power and cooling. And it is this total cost of ownership, not individual drive prices, that ultimately determines the cost-effectiveness of a storage solution.

Improving Storage Efficiency

In the late Nineties enterprise data was largely transactional, so it was only natural to employ costly high-performance drives for virtually all disc-based enterprise storage applications. However, as time passed and the quantity of reference data (including safety copies of mission-critical files) mushroomed, it became clear that more targeted, cost-effective storage was needed.

Fortunately, such high-efficiency storage has now become a reality. Purpose-built disc drives are widely available to ensure every application—from online transactional duties to nearline bulk storage tasks—can be addressed with the optimal blend of performance, reliability and cost-effectiveness.

SAS is leading this movement, its compatibility with SATA ensuring unprecedented freedom to specify and consolidate the most appropriate storage devices for a broad range of applications. Beyond the cost savings of utilizing a common infrastructure, SAS/SATA compatibility promotes further efficiencies by enabling IT managers to:

- Fine-tune storage cost/performance for a given application with the optimal blend of SAS and SATA drives
- Fully exploit SATA's cost/GB economies without compromising enterprise management and reliability requirements
- Increase storage capacity while maintaining application requirements

Online Storage: 15K SAS Is King

Online storage application requirements can be neatly summed up with an age-old aphorism: Time is money. Transactional data is the very lifeblood of many enterprises, and online storage is all about keeping that data flowing. In the online world, diminished throughput and bouts with downtime are far more than inconvenient; they often translate into substantial lost revenues.

The advanced architecture of SAS (point-to-point connectivity, full duplex operation, dual ports) eliminates the significant performance bottleneck imposed by parallel SCSI's shared-bus topology. It makes little sense to deploy a SAS infrastructure and then effectively reintroduce a performance bottleneck by using slower SAS drives. While both 10K SAS and 15K SAS drives offer exemplary reliability in transactional environments, only 15K SAS can boast the highest available throughput for online applications.

Simply put, 3.5-inch 15K SAS disc drives are purpose-built to address the performance deficiencies of 10K SAS drives, and in so doing enable greater cost-effectiveness in storage. Requiring fewer drives to provide equal or better performance, 15K drives deliver lower TCO by ensuring:

- Less supporting infrastructure
- More enclosure/rack space for nearline storage
- Lower maintenance and storage-management costs
- Greater reliability due to fewer potential points of failure
- Higher IOPS/U

15K drives achieve their superior performance by providing quicker access to data. Of course, 15K drives spin their internal discs faster which reduces latency. But they also utilize smaller diameter discs (approximately 2.6 inches vs. 3.0 inches for many 10K drives) to shorten the distance an actuator must move, decreasing seek time.

Costly 10K Kludges

Some IT departments may attempt to meet their online application needs with 10K SAS drives, employing unorthodox techniques to wrest improved performance from them. For example, they might deploy lower-capacity drives to ensure each drive accesses less data, thus lowering drive access times. In effect, these departments would be purchasing more actuators to simultaneously access their data.

Or they may resort to storing data only on the outer diameter of the disc, a practice referred to as *short-stroking* (or *destroking*). This reduces the distance the actuator must move to access the data, thus improving disc seek time at the cost of severely limiting effective capacity.

Both of these techniques vastly increase storage costs due to the high number of drives needed to meet a given capacity requirement, and this also drives up associated costs such as infrastructure, required space (data center ft³ and rack U), power consumption and cooling, and storage maintenance.

Nearline Storage: Nearline SATA vs. 10K

Nearline applications can be divided into two categories: data protection and reference data. Both entail data that doesn't justify the cost of high-availability, mission-critical storage, but must still be readily accessible at any time. While nearline data activity is far less frequent than online activity, both involve random reads/writes that force drive heads to rapidly and repeatedly traverse a drive's discs.

To deliver the nearline-class reliability standard of 1.0 million hours MTBF, nearline-ready SATA drives are specifically designed to withstand the rigors of random reads/writes and 24x7, always-on operation. In contrast, the typical 600,000 hours MTBF rating of desktop-class SATA drives is obtained in the mild environment of sequential reads/writes and 8x5 power-on hours, and thus has no relevance when considering the use of such drives in nearline applications.

Nearline-ready SATA drives also incorporate Workload Management to dynamically protect them from excessive peak workloads. To further safeguard reliability, these drives perform "offline scans" during drive idle time to periodically test the media surface for defects.

Designed for online applications, 10K SAS drives offer even greater performance and reliability than nearline-ready SATA drives. However, when used in less demanding nearline applications these performance and reliability advantages are largely academic. More importantly, 10K SAS drives are far more expensive in terms of capacity, with cost/GB *roughly four times higher* than nearline-ready SATA drives.

Conclusion

Serial Attached SCSI will play a key role as the enterprise continues to move towards more optimized, cost-effective storage platforms. A single SAS-based subsystem equipped with 15K SAS drives and nearline-ready SATA drives is capable of handling the full gamut of enterprise storage duties, from online (transactional, high availability) to nearline (archival, low availability) applications. Purpose-built for nearline duty, this new breed of nearline-ready SATA drives ensures application reliability is not diminished.

Moreover, the peerless performance of 15K SAS drives enables systems that deliver greater throughput with fewer drives than 10K SAS systems, significantly increasing configuration flexibility and cost-effectiveness. Unrivalled in both performance and efficiency, 15K SAS-based solutions exemplify storage without compromise.