

The Rise of SAS as NextGen Data Storage

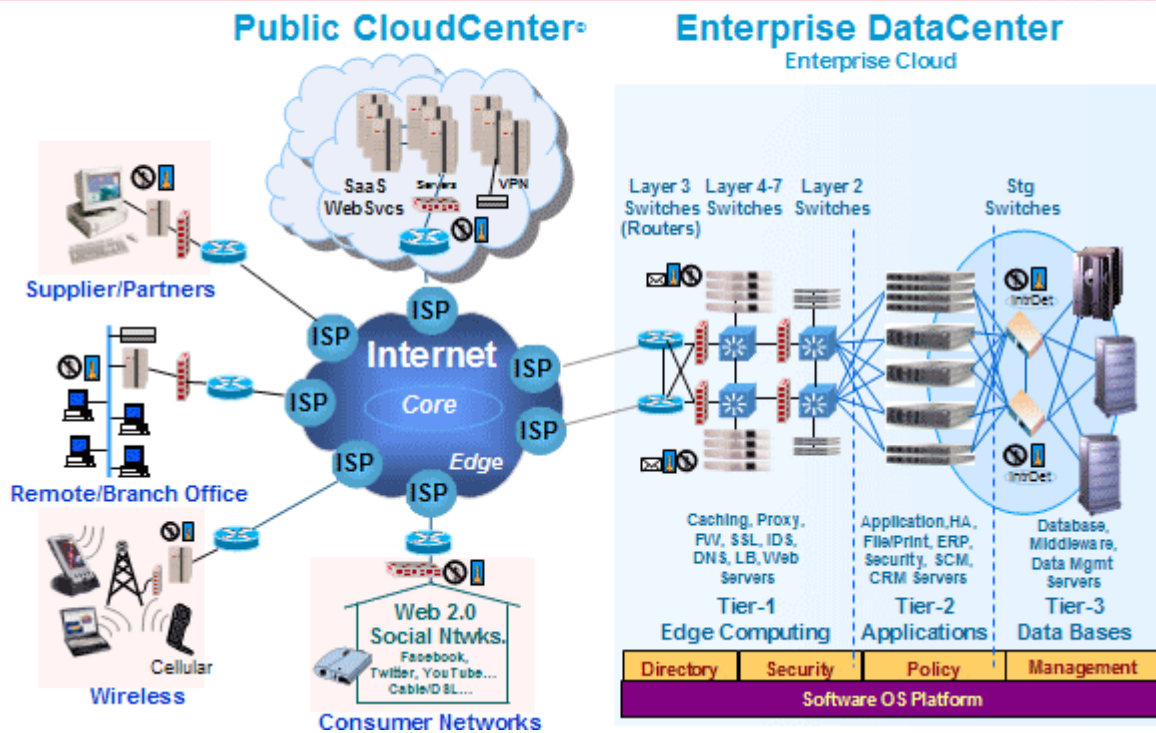
[Anil Vasudeva](#)
Senior Editor
[IT Newswire](#)

Industry Trends

Explosion in the Growth of Storage Data

The internet has been the catalyst behind the explosive growth of digital information, growing at an annual rate of 60% and estimated to reach 1,800 Exabytes by 2011. Regardless of the state of the economy, the amount of data that needs to be stored, accessed, and managed will continue to grow exponentially, especially as new types of data, such as that from social networking and Internet video, continues to mushroom.

► End to End IT Infrastructure with HA & Security



Based on IMEX Research work with the providers and deliverers of Web 2.0 content and services, three fundamental trends stand out that would dictate the requirements of the new infrastructure architecture:

- The infrastructure needs to be designed to scale quickly in order to dynamically react to the fluctuations in demand for capacity, performance, and high availability of stored data.
- Storage has to be managed holistically - with an ability to deploy a 10 PB system and manage the data easily with a small staff.
- With the increasing shift to support Web 2.0 and Cloud Computing and Services, the economics of the infrastructure has to align with the new business model, needing an inexpensive infrastructure and truly cheap storage for this new business model to work. Google Search, YouTube Video, Facebook or Flickr collaborative sites could not have grown if they charged \$10 a month to access them. Their monetizing models have evolved to soft-sell viral ads, instead. In such cases a PB of data at \$15/GB wouldn't sell, nor would \$3/GB if you need an army of administrators to keep the system up and running.
- Enterprises have started to evolve into Enterprise Datacenter/Private Enterprise Cloud (using Service Oriented Architecture (SOA) running Software as a Service (SaaS)) and Public CloudCenter©

(run by Service Providers) and supporting compatible SOA and software wherein the applications can be equivalently run in either the Private Enterprise or the Public CloudCenter.

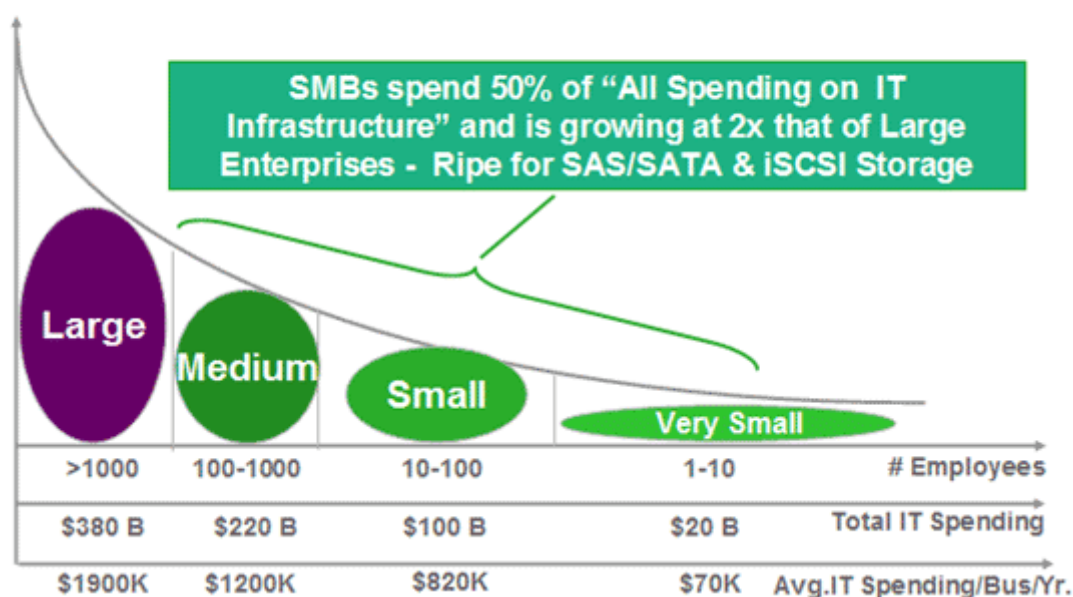
This article describes how the storage industry is providing newer, cost-effective Serial Attached SCSI (SAS)/Serial Advanced Technology Attachment (SATA) storage technologies and platforms to keep pace with this growth and required operational metrics.

Market Drivers for SAS/SATA Storage

SMB Markets - Perfect Storm for SAS/SATA

Sharing data by multiple PC clients over a Local Area Network (LAN) as companies grew, gave rise to LAN-attached storage (NAS). Access to shared storage by clients became a bottleneck whenever backups from disk NAS to tape drives happened, as it completely hogged the LAN bandwidth. That required a separate dedicated network (a la SAN) to do LAN-less backups without disturbing the client's direct access to NAS. These SAN networks had to be lossless, unlike Ethernet based LANs which would wilt under high traffic and start dropping packets as required by their CDMA design, resulting in constant retries and a complete loss of performance. As a result, a lossless (albeit very expensive) Fibre Channel Storage Area Network (FC San) was devised. It needed special components but could be used over long distances as required by large companies with operations spread over multiple divisions, which only their large dedicated storage IT staff could implement and maintain.

► The SMB Long Tail



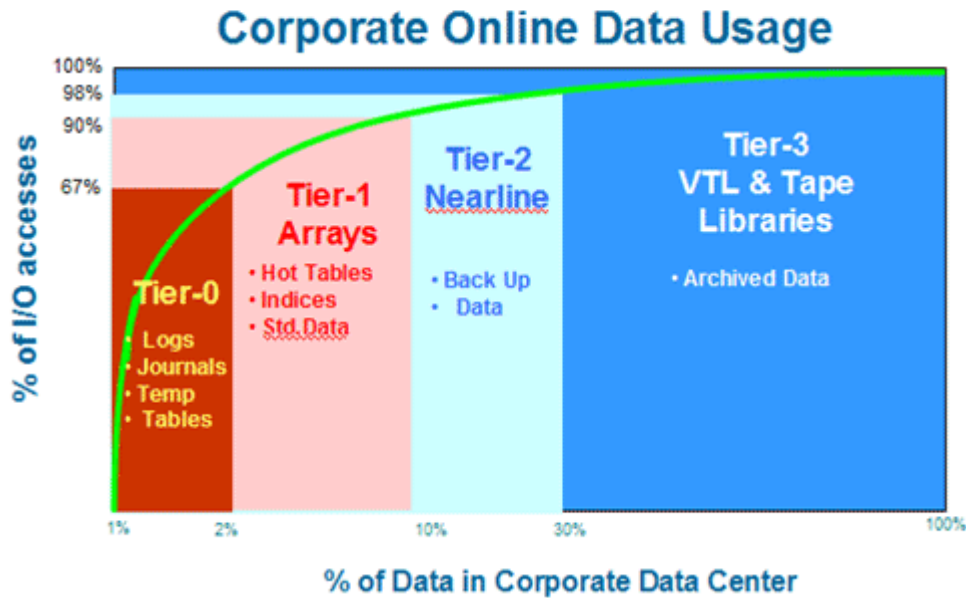
In the case of small-to-medium-size businesses (SMBs), SAS fabrics are being implemented in configurations where expensive FC SANs are not attractive. These are configurations where high performance is needed and the solution needs to be low cost, yet doesn't need to cover large distances such as a small scale datacenter with less than 50 servers and storage arrays. By using a SAS fabric, the data being written to and read from the hard drives does not need to be translated to a different protocol to move thru the fabric as in some FC-based SANs. This creates a perfect storm for matching the needs of SMB markets with SAS/SATA storage solutions. Companies implementing their IT infrastructure using blade server-integrated storage within a small chassis with multiple chassis in a rack, as the "datacenter in a room," are perfect candidates for the use of SAS/SATA storage.

Tiered Storage in the Enterprise

Given the different cost/performance/availability required by different workloads, it is natural to optimize the storage infrastructure to optimally match that need. This resulted in the generation of multiple tiers in storage systems, namely:

- Tier-0 Solid State Flash Drives
- Tier-1 Primary Storage (performance-optimized SAS/SATA based disk arrays)
- Tier-2 Nearline Back up Storage (capacity-optimized SATA storage)
- Tier-3 Archival Storage (SATA-based VTL and Tape Libraries) as shown in the diagram below.

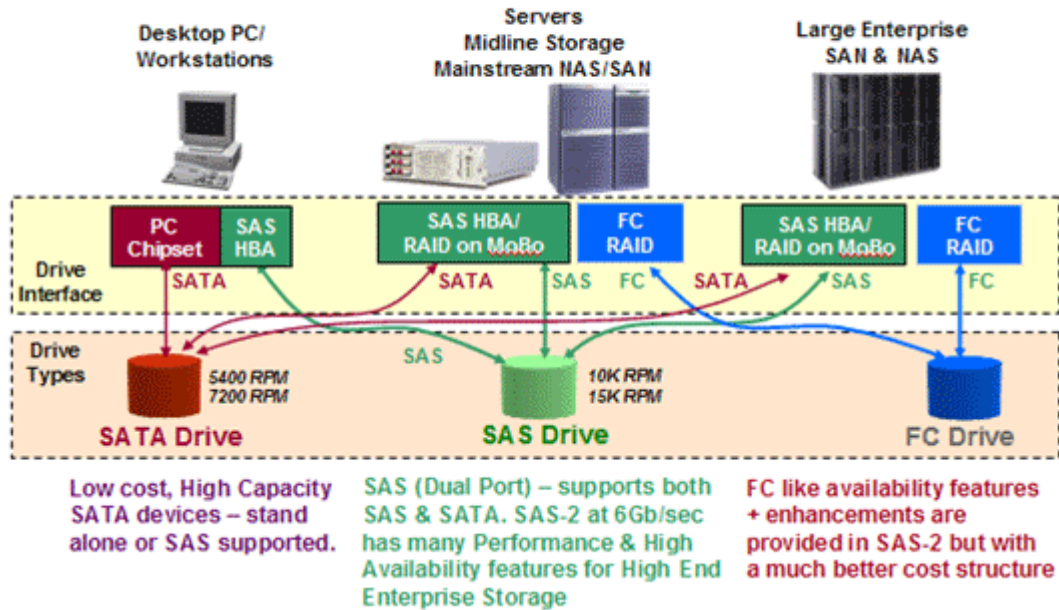
► Corporate Data Usage



SAS Evolves with the Industry

The storage industry responded by delivering high-capacity drives that store 3-4x more data than traditional performance-optimized drives, and are suitable for 24x7 use in rack-mount environments. These high-capacity drives are helping IT managers meet their storage growth requirements and will continue to increasingly penetrate the datacenter as users learn how to better leverage these storage devices. With its ability to support both SAS and SATA disk drives, SAS is also making headway as the disk-drive interface of choice for external storage in both JBODs and external RAID subsystems. SAS is beginning to penetrate a segment that until now has primarily been the domain of Fibre Channel. In 2007, roughly 10.8 million capacity-optimized HDDs shipped in datacenter storage solutions. Today SAS is the dominant storage interface offered by server vendors worldwide.

Tiered Storage by Price/Performance



Enterprise Level Storage System Features SAS is Scalable, Sharable, Secure Storage

Six Gb/s SAS is more than an improved set of features over the current generation of SAS. It offers IT managers highly tangible benefits that will make their data more reliable, secure, and faster. SAS-2 (T10 SAS standard designation) allows up to 10-meter cables, standardized expander zoning and spread spectrum clocking to reduce radiated emissions (EMI). Multiplexing, which allows multiple, slower speed data streams to be aggregated into a 6Gb/s data stream is an efficient way of aggregating bandwidth. SAS-2 provides performance and row address strobe (RAS), which are at least on par with that provided today by Fibre Channel. 6Gb/s SAS products and systems will make a significant impact on storage in 2009.

T10 Committee SAS Standard Designations and Features

	SAS-1 & SAS 1.1	SAS-2	SAS 2.x
Industry Trends	<ul style="list-style-type: none"> Blade Servers Serial Storage High Capacity Drives 	<ul style="list-style-type: none"> Virtualization, Clustering Multi-core Processors Large Topologies 4K Sector Size Shared External Storage 	<ul style="list-style-type: none"> DataCenter Scale-Out Fabric Connectivity Greener Storage Cloud Computing
Key Attributes & Features	<ul style="list-style-type: none"> SCSI Legacy SATA Compatibility 3 Gb/s RAID-6 SFF HPC 	<ul style="list-style-type: none"> 3 Gb/s Compatible Improved Signaling 6Gb/S, Scalability Zoning Management Data-Protection (DIF) Security (FDE) Multiplexing Multiple Links Support for Virtualization, Clustering, SSDs Improved Ext. Storage 	<ul style="list-style-type: none"> Up to 20m distance Lower Power Features 6Gb/s Rack-to-Rack (~20m) MAID Support Power/IOP/Unit Switching

Some of the main features of SAS-2 include:

- Performance — doubles the link rate and bandwidth
- Multiplexing — optimizes bandwidth by enabling two 3Gb/s links to share a 6Gb/s port

- Increased zoning capabilities — enables partitioning of a domain into smaller sets of accessible devices
- Self-configuring expander devices — accelerates system initialization and change detection
- Diagnostics and robustness — improves status reporting and error notification

Flexibility - Combining High Performance with High-Capacity Drives

SATA drives are primarily designed for cost-effective bulk storage. To achieve economies of scale, SATA drives feature lower spindle speeds (typically 7,200 rpm), lower mean-time-between-failure rates, and lower cost. Consequently, they tend to be applied where transaction rates are low and data availability is not critical.

SAS drives, on the other hand, are built for high-performance, high-availability use. SAS-2 has the ability to connect high-capacity disk drives (SATA) alongside performance drives (SAS) in a storage system. The SAS connector itself is designed as a single, uniform backplane, so designing a system with both drive types is simple. This compatibility reduces the cost and complexity of storage designs, since SATA devices are fully compatible with SAS controllers - the SATA Tunneled Protocol (STP), included within SAS, passes SATA commands through to the SATA drives.

SAS and SATA compatibility also allows system builders to design hybrid storage systems using common connectors and cabling. Installing or upgrading either SATA or SAS drives in the same system is simply a matter of replacing one drive type with the other as the SAS backplane connectors receive both SAS and SATA devices. However, since SATA backplanes connect only to SATA devices, backplanes should use SAS connectors to provide the greatest system design flexibility.

SAS connects with SATA through one of the following techniques using:

1. Expanders with SATA Tunneling Protocol (STP)/SATA bridging
2. SATA drive tailgate cards with Serial SCSI Protocol (SSP)/SATA bridging
3. High-capacity SAS drives with pure SSP.

Each approach has its advantages and disadvantages that should be considered during architectural design.

Cable/Connector Consolidation

To meet industry's demand for denser cabling solutions, the small form factor (SFF) mini-SAS connector has been quickly adopted for both internal and external connectivity.

Scalability

Datacenters require storage architecture that is able to scale on demand. By shifting more of the SAS topology discovery process from the host controller to the expander, and by providing the added capability of flexible table-to-table routing, SAS-2 now dramatically reduces SAS messaging during topology discovery, resulting in reduced time to discover, initialize, and scale ever-increasing devices needed by large, tiered-storage solutions.

SAS uses expander hardware to simplify configuration of large external storage systems that can be easily scaled with minimal latency while preserving bandwidth for increased workloads. The expander hardware enables highly flexible storage topologies of up to 256 mixed SAS and SATA drives. SAS expander hardware, in effect, functions as a switch to simplify configuration of large systems that can be scaled with minimal latency degradation while preserving bandwidth for increased workloads.

SAS today has the ability to connect multiple servers and thousands of storage devices. Scalability at that level often requires that the storage devices and/or subsystems be consistently assigned, or zoned, to operate with multiple hosts in virtualized server deployments. This ability to assign various operating domains for both shared and separate pools of storage is accomplished through a capability scheme referred to as *SAS expander zoning*. This standardized zoning improves SAS' ability to effectively support more complex topologies across multiple expander vendors, while increasing the number of supported zones from 128 to 256.

Intelligent Self-Configuring Expanders

Expanders are capable of implementing self-configuration features. Each expander device discovers the devices attached to it and completes its own route table. Since all expanders are initializing at the same time, the overall system topology is resolved quickly.

SAS as a Fabric

SAS was developed as the natural evolution of parallel SCSI, enabling point-to-point drive connections via a serial interface. To support direct-attach storage outside of the server, the concept of an expander was defined. SAS expanders enable a simple switching topology and allow multiple servers to connect to the same SAS JBOD, and then be shared between multiple servers.

As larger and larger SAS-based topologies are being implemented, there is discussion over using SAS as a fabric technology. To fully understand this phenomenon, it's important to understand the roots of SAS and how SAS systems are being architected. SAS-2 provides additional status and reporting information to facilitate diagnostic functions. In the event of a fault, this status data can be used to identify, isolate, and analyze fault and error conditions.

SAS expanders, built into SAS switches, enable SAS fabrics to be built. SAS fabrics make it simple to add more storage to a configuration. SAS hard drives in a SAS fabric are sharable to all servers connected to the fabric. SAS zoning allows administrators to divide the storage into segments and then decide which servers are allowed to access which specific segments. Server blades connecting to multiple storage blades form a good environment for use of a SAS fabric.

Command Queuing

Like SCSI, SAS includes advanced command queuing with 256 queue levels, providing unique intelligent data handling features such as head-of-queue and out-of-order queuing. These queuing features are critical to enterprise applications because they allow a system to reorder and reprioritize commands within the interface.

Increased Security through Zoning

Secure zoning is a part of the new SAS-2 specification which is defined to increase security in a server and storage environment. A "zone" is similar to a hardware firewall that compartmentalizes a group of disk drives to create secure zones, segregating data. It works with 6Gb/s and 3Gb/s SAS and SATA disk drives within a 6Gb/s SAS environment.

SAS-2 zoning enhances the SAS fabric by providing a hardware mechanism to increase device segregation. SAS data storage systems may include a variety of device types such as SAS and SATA, as well as data protection mechanisms such as RAID and encryption. Zoning enables segregation of these storage types at the system level to simplify partitioning, provisioning, and overall system management. Zoning can be optionally secured by password to prevent unauthorized access, malicious attacks and corruption of data by operator or application error on the server.

Enterprise-level Data Integrity

Commonly referred to as Data Integrity Field (DIF), it allows both data and commands to be protected from the application layer, all the way from the host to the storage system to the disk drive.

Supporting Industry Standards-based Ecosystem

Decision Feedback Equalization (DFE) allows SAS cabling of up to 10m at 6Gb/s transfer rates, keeping pace with the throughput being offered in PCI Express 2.0 servers. At 6Gb/s, the second generation (SAS-2) with 6Gb/s controllers are optimized to take full advantage of the 5Gb/s per-link speeds of PCIe 2.0 enabling peerless system robustness. The improvement in bandwidth allows more disk drives to be added to the high-performance SAS links without the need for additional host controllers or ports, freeing up PCI Express slots for other system expansion needs and reducing cable congestion.

Faster Performance

Network user demands for faster data keep growing. SAS-2 helps overcome the increased demands by doubling the throughput capability of SAS to 6Gb/s. Each SAS connection now supports up to 600 MB/s of throughput. Common SAS controllers come with four or eight ports, creating connections up to 2.4Gb/s and 4.8Gb/s of throughput, respectively.

The full-duplex, point-to-point nature of SAS enables simultaneously active connections among multiple initiators and high-performance SAS targets. Narrow ports allow for a single serial link, while wide ports support multiple links, allowing the aggregation of eight SAS or SATA targets to increase total available bandwidth to 24 Gb/s, the significant bandwidth requirement of large SAS topologies.

Moving to 6Gb/s SAS means faster data throughput, yet it works with 3Gb/s SAS, which protects any current investment in SAS disk drives and storage systems. With the use of 6Gb/s SAS expanders, twice as many 3Gb/s disk drives can be connected with the 6Gb/s SAS multiplexing capability.

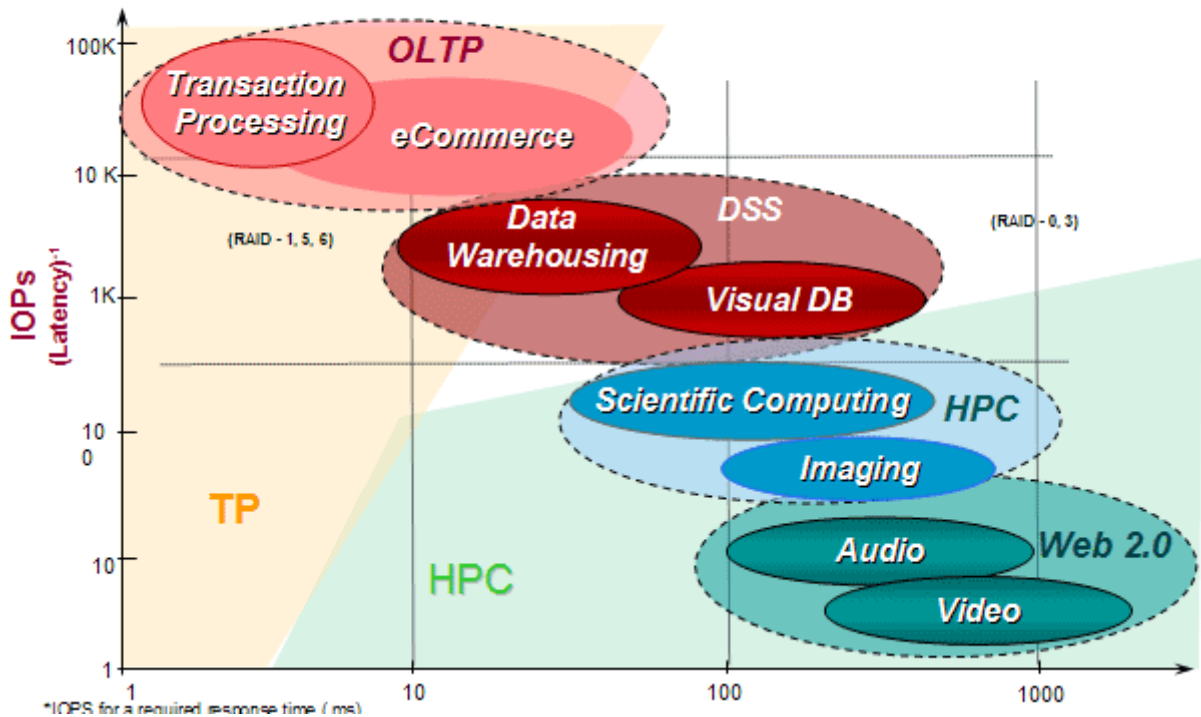
Solid State Drives to the Rescue

Solid state drives (SSDs), have a promising future in the enterprise space. They promise to overcome literally all limitations of traditional hard drives – power consumption, heat dissipation, mean-time-between-failures, speed and I/Os, etc. Much as 2.5" SAS HDDs can now offer high performance (at both 10,000 RPM and 15,000 RPM spindle speeds) when used in conjunction with SSDs, they can significantly improve overall system IOPs – typically about 30 times – with response times of less than two milliseconds as compared to a 15,000 RPM Fibre Channel drive. Commercial databases can get high levels of value from these when they are used intelligently, which is to say sites can put metadata rather than data on them. SSDs represent what may come to be thought of as a new storage tier—*tier 0*.

For all of their good features, some of the shortcomings of SSDs relate to their wear characteristics resulting in limited read cycles. To overcome such limitations, SSD vendors have created wear-leveling techniques wherein failing or bad data blocks are diagnosed early on and automatically removed and substituted under a controller, thus mitigating reliability concerns. In fact, the larger the SSD storage size, the better are these wear-leveling algorithms perform. NAND-based SSDs have a growing opportunity in the datacenter. Similar to HDDs, flash-based SSDs are offered with several interface options. Most manufacturers are standardizing SSDs with a SAS/SATA interface instead.

Application-Aware Storage Infrastructure

► **Market Segments by Workloads**



Certain storage applications requiring high IOPs to support the data architecture and performance requirements of the computer system are candidates for SSD-enabled acceleration. (See chart: Application-aware Storage Infrastructure) above. Enterprise storage applications can strongly benefit from the use of SSDs in conjunction with cost-effective SAS/SATA drives.

The use of SSDs, besides improving performance, can also significantly reduce power consumption since they have no spinning media. SSDs consume a fraction of the power consumed by magnetic hard drives. A 64GB flash drive, for example, can use 30 to 40 percent less energy than a 73GB 15,000 RPM magnetic drive. Reduced power consumption also means reduced heat dissipation. As such, the array as a whole will have a lower thermal footprint and reduce air-conditioning requirements.

Investment Protection

SAS-2 will have a 6Gb/s data rate and be backward compatible with existing 1.5Gb/s SATA and 3Gb/s SAS/SATA products and infrastructure. The 6Gb/s SAS interface not only enables faster data rates, but also offers new benefits and opportunities for enterprise applications, including the ability to spread I/O requests over a greater number of HDDs. It also provides a higher performance interface for future SSD devices designed with very fast I/O and data-rate capabilities. The 6Gb/s SAS interface also allows for the design of SAS storage solutions that could compete with storage systems currently leveraging the Fibre Channel interface.

Backward Compatibility

SAS-2 at 6Gb/s doubles the previous generation's bandwidth for each link and adds link multiplexing to enable a 6Gb/s link rate to share two 3Gb/s connections. 6Gb/s SAS is built to be backward compatible to 3Gb/s and 1.5Gb/s link rates. The SAS infrastructure supports a mix of SAS and SATA drives and link rates.

Scaling SAS Implementations

One of the benefits of SAS expanders is that they can be cascaded, enabling very large configurations to be built. A single SAS domain may support up to 16,384 devices and access every hard drive. The SAS-1.1 specification enables the concept of SAS zoning, whereby a configuration of SAS hard drives is broken into groups or zones and servers are enabled to communicate with drives in one or more zones. IT managers are able to specifically prevent some servers from communicating with some zones. SAS zoning is implemented in the SAS expander and this new expander essentially becomes a SAS switch. Once SAS switches are incorporated into a configuration, a SAS fabric is built. Such a switch is much less expensive to build and maintain than Fibre Channel. SAS fills the need for a high-performance, low-cost fabric that is not required to span very long distances.

SAS and SATA Compatibility

One of the main reasons that SAS has been able to scale is due to its compatibility with SATA HDDs, which provide the highest capacity at the lowest cost-per-gigabyte of any storage media. In addition, the use of a SATA Active/Active port selector to dual-port a SATA HDD enables fully redundant storage architectures with greater system fault tolerance. Where SATA drives are used for infrequently accessed data, such as near-line storage or backup, a Redundant Array of Independent Disks (RAID) is commonly used to mitigate the reliability risks of SATA storage.

The ability to support enterprise quality data storage with SAS, and cost-effective, high-capacity storage with SATA, both using the same SAS infrastructure, has led to economical, scalable storage and server offerings. In addition, having a mix of SAS and SATA drives allows for information lifecycle management (ILM), whereby data migrates from primary 24/7 storage using SAS devices, to secondary/nearline storage using SATA devices as it ages and is accessed less frequently. When data has completed its useful life, it is moved to tape for archiving.

Management

The Serial Management Protocol (SMP) is enhanced to provide more configuration, faster initialization, and greater reporting for diagnostic and status monitoring.

SAS and Windows

SAS' backward compatibility with previous-generation SCSI software and middleware makes it easy to incorporate legacy components—hosts and drives—into evolving SAS topologies, eliminating new training or integration costs and the need for modifications to legacy software.

SAS Market Penetration

In addition to the rapid transition to SAS for performance-optimized drives, the vast majority of capacity-optimized drives shipping for enterprise applications today employ the SATA interface. When accounting for all performance and capacity-optimized HDDs that shipped in 2009, more than 70% shipped with a serial interface (SAS or SATA).

SAS leverages proven SCSI functionality and builds on the enterprise expertise of multiple chip, board, drive, subsystem and server manufacturers throughout the industry. In the enterprise, SAS has crossed over to being pervasive in the industry.

The Future - Steps for SAS 2.X

In order for SAS to keep pace with the ever-increasing needs for more capacity and more complex capabilities beyond 6Gb/s SAS, additional enhancements are planned for the 2011 timeframe—currently referred to as SAS 2.X. The main focus of improvements includes:

- Datacenter Scale-out Capabilities – providing improved cabling options – copper cables of 20 meters or more, and the potential for optical connections for even longer cabling distances.
- Energy-efficiency Green Storage Features – providing power management options that would bring SATA style power management into the SAS system to improve power and cooling efficiencies.

When these enhancements are added to 6Gb/s SAS—providing even longer distances, larger infrastructure support, and improved power management capabilities—they will allow for larger data-system scale-outs and also generally greener storage.

High Capacity SAS/SATA to Mitigate Data Proliferation

With virtually every IT department in today's corporate world facing growing user demands and shrinking budgets, storage vendors are rushing to deliver the cost efficiency of SAS/SATA systems with value-added features and availability levels typically found only in enterprise class facilities.

Energy Efficiency

The best path to storage efficiency is probably the use of data deduplication. Every major vendor now offers this capability (sometimes also referred to as "single instancing"), which eliminates redundant data, replacing redundant bytes or blocks with pointers, which take up much less space. Deduping can reduce the overall need for storage capacity by 10%-30%, depending on the application and the kind of data being stored. Deduping is an easy way to save on storage costs and on the power needed to drive arrays, and it doesn't necessarily require committing to new architectures.

MAID (Massive Arrays of Idle Disks) systems spin down their disks until the data on them is needed, at which point they spin up again for as long as they are in use. MAID is finding a home as a storage medium for "persistent" data, information that needs to be available for reference, but for which speedy access is not necessary.

Opportunities for Embedding Intelligence in SAS Expanders

Not only is computing infrastructure performance important but exponential demands for new storage facilities have lately been driven by legislative requirements including Sarbanes-Oxley, HIPAA and others. The costs of storage technology acquisition (capital expense) have fallen below that of ongoing operational and maintenance (operational expense) of exploding data storage.

New opportunities exist in embedding intelligence for deduplication, MAID, and workload-aware dynamic provisioning of virtualized storage at the SAS expander level.

Towards a Scalable NextGen Datacenter (NGDC)

The path to a NGDC is through SIVA—Standardization, Integrationm Virtualization, and Automation, as shown below.

▶ Follow SIVA® for a Scalable & Dynamic NGDC

