



Breaking the I/O Paradigm — SAS Enters the Nearline Storage Race

Analyst: David Reine

Management Summary

Over the past century, men and women have not only become increasingly dependent on their automobiles, we have also fallen in love with them. They have become a treasured part of our daily routines, commuting to work, shopping, or visiting relatives on the holidays. They have become more than just a mode of transportation. We use them to reflect our personalities: free-spirited convertibles, high-strung muscle cars, and practical family sedans. Many times, however, our budgets do not always match our image. We must strike a compromise between cost and performance, between appearance and practicality. You cannot fit a family of four and a dog into a sports car nor can most of us afford to fill up the gas tank for a 350-horsepower engine that can go from 0 to 60 (or 55) in six seconds.

Over the past few decades, however, technology has come to the rescue, improving the performance and appearance of relatively inexpensive vehicles, while maintaining a relatively low total cost of ownership with 5-year warranties and improved mileage economy. One of the more intriguing developments of late is the introduction of hybrid engines, capable of running on electricity part of the time, conserving the amount of gasoline required, preserving a natural resource. Early hybrid cars enabled spectacular fuel economy, but employed weird shapes and had few luxury options. Newer hybrids use their gas-electric engines for more acceleration and power. Automakers such as Toyota, Lexus, and General Motors are building hybrids with the look and feel of regular cars. There has been a compromise between performance and economy, between luxury and the environment.

A similar evolution is occurring in the Information Technology (IT) arena. Storage architecture continues a transition that began with direct-attached SCSI devices in the 1980s and evolved to the high-speed storage area networks (SANs) of the 1990s, with expensive Fibre Channel drives and specialized technical support. Since 2003, we have seen an influx of lower-performing, but inexpensive SATA drives. Now we see the introduction of low-cost, high-performance Serial Attached SCSI (SAS) to change the storage I/O paradigm further. We do not have to compromise between cost and performance. Now the data center can reduce the total cost of server ownership by implementing a multi-tier architecture that enables information lifecycle management, or ILM, within a single cabinet, as we see SAS backplanes supporting both SAS and SATA drives, enabling a storage solution upgrade without the forklift. To learn more about SAS, please read on.

IN THIS ISSUE

- **Data Center Storage Requirements 2**
- **The Role of SAS in the Enterprise..... 3**
- **Conclusion 4**

Data Center Storage Requirements

Over the past decade, we have seen an explosion in the demand for additional storage within the data center. In an attempt to get control of the situation and enable a sound storage management strategy, CIO's and IT managers have implemented a policy of consolidation in order to keep the total cost of ownership (TCO) for storage within the constraints of already reduced budgets. One of the first lessons learned was that all data is not the same: some data is more valuable than other and some data changes in value over time. Storage, in fact, has become a three-headed monster, requiring different architectures, depending upon which characteristics an application demanded:

- Highest performance,
- Highest reliability capabilities for mission-critical applications,
- High capacity for off-line archiving purposes, and
- A new class of storage – *nearline* – demanding both high performance and high capacity for solutions, such as backup and recovery, but without some of the reliability requirements of mission-critical data.

This last category has been implemented in support of regulatory compliance requirements, especially in the area of financial and health records. In this category, all data has become critical, and has to be retained in a readily accessible manner for years. Enterprises have to be able to manage, **in a cost-effective manner**, the different kinds of storage that are doubling capacity annually.

To ensure that the most valuable data is stored on the fastest and most reliable storage devices, a new management strategy was needed to control acquisition costs and lower the TCO of enterprise IT. Information Lifecycle Management (ILM) was devised in order to minimize the cost to deploy, maintain, and administer a consolidated storage network consisting of multiple tiers of storage devices in order to support *online* (mission-critical), *nearline* (secondary), and *archive* (offline) data. In the course of implementing ILM, the data center has seen a transition in the way that servers communicate with their storage devices. Following Moore's Law¹, the server has seen dramatic improvements not only in processor speed, but also in the size and

speed of memory. The capacity of disk drives has also increased rapidly over the past two decades. Unfortunately, the improvements in performance of disk I/O controllers pale in comparison.

For the past twenty years, parallel SCSI has been the principal storage communication protocol in the data center, with ATA used as a low cost alternative. Unfortunately, parallel SCSI, introduced in 1986 at 2 MB/s, reached its ceiling with the Ultra-320 implementation for up to 15 devices. The requirement for rapid access to critical information from multiple servers to multiple storage devices, both local and remote, with minimal latency, led to the development of Fibre Channel (FC) in 1995. This was the first serial technology for open systems storage to gain acceptance in the enterprise data center². FC was introduced at 1 Gbps, but is now available at speeds up to 4 Gbps. Unfortunately, the rapid access to mission-critical, or primary, data comes at a high cost, thus becoming unacceptable as a storage medium for secondary or archive data. In order to control the costs for high-capacity storage requirements where mission-critical performance and reliability are not required, *Serial ATA (SATA)* was developed. SATA I replaced the low-cost, high-volume desktop parallel ATA interface³, eliminating limitations caused by:

- A short, bulky ribbon cable that impeded airflow and was difficult to install;
- Multiple wires and signals that were prone to electrical interference; and
- The lack of a hot-plug capability.

SATA serialized the interface at 150 MB/s, improving upon the 133 MB/s maximum throughput of ATA, evolving it into an enterprise role as secondary storage for applications, such as video editing and scientific computing. Despite its lower performance and functionality, low cost and high capacity made it an attractive alternative to SCSI and FC. SATA I replaced ATA in 2004 and has been the alternative of choice to FC at the low end. SATA drives are available from all the major disk drive vendors in a variety of capacities, from 40GB to 500GB, with 2MB to 16MB of cache. SATA I drives can be installed in a hot swap, hot plug mode with point-to-point connections using a low-voltage differential interface to improve the reliability and integrity of the signal, doubling the cable length to 1m. SATA has benefited from a groundswell of support up to now, with adoption by wide array of server and storage vendors.

With ATA drives capable of 60 MB/s through-

¹ Moore's Law originally predicted that the number of transistors on an integrated circuit would double every year and would continue for the foreseeable future. Even though that may have slowed, we continue to see that data density has doubled approximately every 18 months, and this is the current definition of Moore's Law.

² IBM introduced Serial Storage Architecture (SSA) in 1996.

³ ATA replaced IDE as a standard in 1994.

put, a 1.5 Gbps interface would appear sufficient. However, as you aggregate multiple drives through a port multiplier, requirements rapidly increase to where four drives in an array would consume 240 MB of throughput. At this point, a 3 Gbps, or 300 MB/s, interface is required⁴. With the advent of perpendicular recording on disk media, higher densities will lead to terabyte devices and even faster throughput demands.

Signal and cable length constraints created hurdles that parallel SCSI could no longer span. The result of a joint industry effort to replace it has produced a serial successor. Over the past three years, *Serial Attached SCSI (SAS)* has been promoted by the SAS Working Group⁵ as the next step in the evolution of the SCSI standard, providing all of the benefits of Ultra320 SCSI: availability, scalability, manageability, and performance. With speeds up to 30 times faster than parallel SCSI, it provides the headroom required for tomorrow's solutions. **SAS is a hybrid architecture: retaining compatibility with the legacy SCSI command set which has become the standard for enterprise computing, frame format from FC, and physical characteristics from SATA.** As the first SAS products become available, let's take a look at how SAS can ease the burden in your data center and lower the TCO for your enterprise.

The Role of SAS in the Enterprise

Critical, as well as mission-critical, enterprise data continues to grow at a rate that doubles data center storage requirements annually. The availability of a faster, more economical, and more flexible method for accessing and delivering multiple tiers of storage to the server becomes more urgent.

FC will continue to be the option of choice for mission-critical data within the enterprise, especially where optical support or fabric features are required. However, the cost of acquisition and implementation of FC necessitates an alternative for storage that does not require the same performance and resiliency. One alternative might be the adoption of *FATA – Fibre Attached Technology Adapted – drives*⁶, or other low-cost, FC-attached, ATA-class

Exhibit 1 – Drive Features

	<u>RPM</u>	<u>Duty Cycle</u>	<u>MTBF (Hrs.)</u>	<u>Seek Time</u>
SATA	7.2K	Low	1.0M	8.5ms
SAS	15K	High	1.4M	3.6ms
SCSI	15K	High	1.4M	3.5ms
FC	15K	Highest	1.4M	3.5ms

drives. For those enterprises that have already deployed a FC SAN, FATA provides a low cost/GB, high-capacity drive alternative without adding another protocol to the storage administration agenda or having to procure another storage array.

SAS provides an outstanding compromise with lower costs, high performance, and flexible configurations. While not as performant and reliable as FC, SAS delivers SCSI compatibility at serial speeds, outperforming comparable Ultra320 drives, with advanced command queuing and advanced verification and error correction. These features enable SAS to replace SCSI as the dominant interface for DAS, NAS, and SAN environments. It provides significantly better performance and resiliency than SATA. (See Exhibit 1, above, for drive features.) **Furthermore, and most importantly, SAS changes the storage paradigm. SAS adapters are fully compatible with SATA drives, providing flexibility and functionality to the data center that has never been available for enterprises. SAS enables customization of an ILM environment based upon enterprise cost and performance requirements. A system can auto-negotiate with both SAS and SATA drives – upgrades can be accomplished without forklifts or system downtime required. As a result, the data center achieves SATA cost with the point-to-point performance and reliability of SAS, with RAID 6 as an option as a hedge against SATA failure.** SAS may be used as reliable, performant primary storage for the small and medium enterprise or a department within a larger business, as well as secondary storage for the larger enterprise. SATA I is cheaper than SAS, with lower reliability and availability, and remains ideal for long-term archiving.

SAS has better functionality than SATA I for Tier-1 applications or high-availability environments. It provides a full duplex point-to-point architecture for better fault isolation, with multi-path, enterprise-class connections when configured with SAS expanders. Each path can transmit data at up to 3Gbps, reliably, with the current architecture, and up to 1,200 MB/s after 2010, based upon current roadmap. Dual-port drives enable both high perfor-

⁴ SATA II is now becoming available at 3.0 Gbps, still less performant than Fibre, but competitive. SATA III at 6.0 Gbps has already been defined for future release.

⁵ The SAS Working Group was founded by Compaq, IBM, LSI Logic, Maxtor, and Seagate in 2001 to ensure the technology as an industry standard.

⁶ See **The Clipper Group Navigator** dated April 8, 2004, entitled *HP Adapts StorageWorks with FATA for Multi-Tier ILM Storage Environment*, at <http://www.clipper.com/research/TCG2004034.pdf>.

mance and high availability with a failover capability, positioning SAS as a low cost alternative to FC in the data center. This is especially true for near-mission-critical, transaction-oriented applications where random data access and data redundancy are critical, such as on-line sales and bank transactions. With the installation of expanders and thinner serial cabling⁷, up to eight meters in length, to improve airflow and routing inside the server, the data center can configure up to 128 SAS and SATA drives on a single edge expander with up to 16,384 drives in a SAS domain through a single fan-out expander. These drives can be low-cost 3.5" 10k RPM drives or high-performance 3.5" drives running at 15k RPM. **SAS provides the flexibility, scalability, and performance unattainable with SCSI or SATA.**

The constantly evolving nature of the SAS architecture also enables a sharp decrease in the TCO of the data center through the use of 10k rpm 2.5" SAS drives. This new development enables a higher degree of storage consolidation than ever before with eighteen 2.5" drives occupying the same space as ten 3.5" drives in a typical array. These small form factor (SFF) disks provide a high-density platform for space and power efficiency, consuming 70% less space and 40% less power than comparable 3.5" devices, enabling better airflow, and reducing the amount of cooling required. In fact, the data center can deploy a 1U rack-mounted server with six hot-plug SFF drives with full RAID 5 functionality in half the space required by a 2U server configured with 3.5" drives, and also enables redundant storage configurations in dense server form factors such as blade servers. This has led to the creation of a new measurement category, *IOPS/U*⁸. A full SAS hardware infrastructure is now available with drives, expanders (multi-port switches to expand bandwidth), and host bus adapters, available in both PCI-X and PCI Express configurations. 146GB SFF drives are now available, with 15K RPM drives to follow in 2007.

SAS also contributes to a reduction in TCO from a software standpoint by protecting the enterprise investment in legacy storage management and application solutions through SCSI compatibility. No workarounds or new driver solutions are required. Additional SAS features include the capability for native command queuing and support of LUNs greater than 2TB.

The SAS initiator supports four standard SCSI protocols, as follows.

- **Serial SCSI Protocol (SSP)** – for full duplex connections, increasing potential throughput to 6Gbps;
- **Serial Management Protocol (SMP)** – to monitor connections between individual devices and to identify points of failure or reduced performance;
- **Serial ATA Tunneling Protocol (STP)** – for connection of SATA devices through expanders; and
- **Half-duplex connections** – for SATA devices connected to SAS controller.

Conclusion

Enterprise needs have outgrown the parallel legacy protocols that drive existing data center storage. CIOs have been forced to make compromises with regard to replacing their storage interface: cost versus functionality, capacity versus performance. In fact, most enterprises have both transaction and reference data. They need to be able to have the high performance for frequent access for transactions, and high capacity with low cost for infrequently-accessed information. If your enterprise has already invested in a FC SAN, then the addition of low-cost FC-attached drives may be the most effective alternative from a TCO basis. However, if you have not made that investment, then SAS mitigates those decisions, combining the best attributes of mainstream SCSI, desktop-class SATA, and enterprise-level FC SAN, meeting the needs of mission-critical enterprise applications and providing an investment protection option with lower TCO, with:

- Enterprise class robustness;
- Choice between SAS and SATA, with 2.5" and 3.5" drives;
- Flexible server and storage deployment and upgrade;
- A simplified management solution; and
- FC-like scalability, availability, and performance.

SAS is on the way to becoming the dominant standard in storage networks. It will become ubiquitous and simplify ILM in the data center and lower the TCO for all server environments, scale-up and scale-out, with a high-performance, enterprise storage solution, enabling an increase in profitability for all who choose to join the move.



⁷ Internal SCSI ribbon cables consist of 34 twisted pairs for a total of 68 conductors, with terminators required.

⁸ Input Output per Second per Unit of vertical rack space.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *“Navigating Information Technology Horizons”*, and *“teraproductivity”* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.