

## **Making the Switch to Shared 6Gbs SAS Storage**

Balancing Performance, Availability, Capacity, Energy (PACE) and Economic Effectiveness

**By Greg Schulz**

Founder and Senior Advisory Consultant, the Server and StorageIO Group (StorageIO)  
Author "The Green and Virtual Data Center" (CRC) and  
"Resilient Storage Networks: Designing Flexible Scalable Data Infrastructures" (Elsevier)  
Coming summer 2011 "Cloud and Virtual Data Storage Networking" (CRC)



**January 18, 2011**

**This Industry Trends and Perspectives Paper is compliments of:**

**LSI**  <sup>®</sup>  
**www.lsi.com/sas**

## **Introduction**

Serial Attached SCSI (SAS) is known as an interface for connecting hard disk drives (HDD) to servers and storage systems; it is also widely used for attaching storage systems to servers. This industry trends and perspective white paper describes how SAS is being used by small medium businesses (SMB) and by data centers for high performance enterprise storage connectivity. SAS has become a viable interconnect along side of Fibre Channel (FC), iSCSI, and InfiniBand (IB) for consolidating multiple servers to access and share storage. With the introduction of SAS switch products, SAS enables the connectivity scaling, functionality and performance expected of higher priced Storage Area Network (SAN) options.

SAS provides a cost effective solution to meet performance, availability, capacity, energy (PACE) and economic requirements while enabling more data to be processed, moved, stored and shared in a given footprint density. Footprint density includes available budget (both capital and operating), power and cooling (Green IT), physical floor or cabinet space, along with management tools and support personnel capabilities. Shared direct-attached and switched SAS storage solutions are increasingly being deployed in diverse environments in place of traditional enterprise protocols such as Fibre Channel (FC) or 10Gb Ethernet (GbE) iSCSI Storage Area Network (SAN). In addition, shared and switched SAS storage solutions are being deployed for high performance external storage in price sensitive environments that previously relied on either dedicated direct attached storage (DAS) or 1GbE iSCSI-based solutions.

## **Common data storage challenges**

There is no such thing as a data or information recession. The resulting impact is the need to process, move and safely store more data in a timely, cost effective manner, and a trend that shows no sign of slowing. Common initiatives associated with addressing the demands and challenges of increased data needs include cloud (public and private), virtualization, and consolidation. Applications or business functions driving data growth include financial, medical and health care, life science, energy exploration, digital modeling, simulation, defense, Internet and social networking, media and entertainment, gaming and video security along with traditional back office functions.

## **Background**

Data storage systems, like computer servers, continue to evolve in terms of functionality, flexibility, performance, available capacity, energy efficiency along with configuration options. With the advent of open systems computing there has been a continued move towards standardized and converged I/O and storage interface protocols.

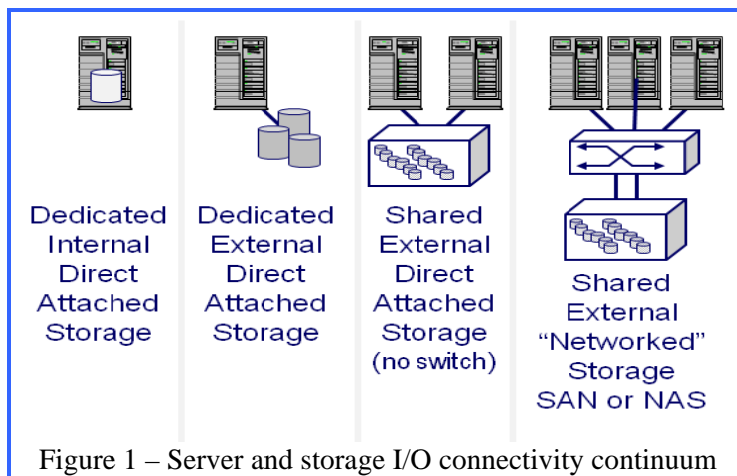
Various storage options for servers include shared external DAS, networked SAN (iSCSI, FC and now SAS) or network attached storage (NAS such as NFS and Windows CIFS file sharing). In some cases, storage is moving offsite utilizing public or private clouds and managed service provided (MSP) capabilities. Over the past decade, networked storage, both SAN and NAS (file sharing), has become more common, but there still remains a large install base of direct attached external and dedicated internal storage. With the diversity of environments and application needs, it is important to have options for resilient, flexible and scalable server, storage and I/O infrastructures. Also, it is important to keep in mind that DAS does not have to mean dedicated internal storage; it can also mean external shared direct accessible storage using SAS, iSCSI, InfiniBand or Fibre Channel in a point to point topology configuration.

## Server and storage I/O connectivity past and present

With the movement toward denser storage systems, bulky cabling with limited connectivity expansion capabilities and distance was not going to scale. In these enterprise or larger scale environments, Fibre Channel has evolved to be a popular option for both server to storage system and storage system to HDD attachment. iSCSI (Internet SCSI) is another popular server to storage system SAN connectivity option where the SCSI command set is mapped to the TCP/IP protocol deployed on Ethernet networks.

Commonly deployed server and storage I/O access scenarios include dedicated internal direct attached storage (DAS), dedicated external DAS, shared external DAS, shared external networked (SAN or NAS) storage and cloud accessible storage. DAS is also called point-to-point in which a server attaches directly to storage systems' adapter ports using iSCSI, Fibre Channel or SAS without a switch.

From left to right in Figure 1 is a server with internal dedicated direct attached storage (DAS), next a server with external dedicated DAS, then two or more servers with shared external DAS.



The DAS I/O or storage interconnect can be point to point using iSCSI, Fibre Channel or FCoE as well as SAS. Continuing from left to right in Figure 1 are multiple servers that connect via a switch to the shared external SAN storage. For example, a SAS, iSCSI, Fibre Channel, Fibre Channel over Ethernet (FCoE) or InfiniBand switch can be used to connect multiple servers to storage devices.

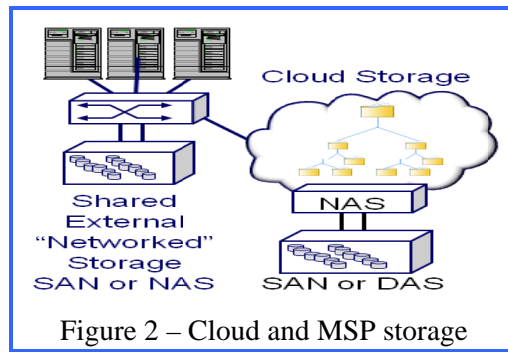
The value proposition, or benefit, of Fibre Channel has been the ability to scale performance, availability, capacity or connectivity over longer distances (up to 10km<sup>1</sup> natively with long range optics) with speeds currently at 8Gbs and 16Gbs on the radar. A challenge of Fibre Channel has been the cost and complexity, which for larger environments is absorbed as part of scaling; however, it is a challenge for smaller environments.

The benefit of iSCSI (SCSI mapped onto TCP/IP) has been the low cost of using built-in 1GbE network interface cards/chips (NICs) and standard Ethernet switches combined with iSCSI initiator software. In addition to low cost for 1GbE based iSCSI other benefits include ease-of-use and scalability. A challenge of iSCSI is less performance compared to faster dedicated I/O connectivity and when a shared Ethernet network is used, increased traffic can impact the performance of other applications. iSCSI can operate over 10GbE based networks, however, that approach requires expensive adapter cards, new cabling, optic transceivers and switch ports that increase the cost of shared storage solution requiring high performance.

<sup>1</sup> Refer to chapter 5 and 6 of Resilient Storage Networks - Designing Flexible Scalable Data Infrastructures (Elsevier) by Greg Schulz

**Making the Switch to Shared SAS Storage**

Figure 2 shows servers accessing shared external networked storage along with remote “cloud” or managed service provider (MSP) based storage services. The cloud based storage service can “behind the scenes” use a point to point configuration shared SAS or a switch topology with SAS, iSCSI, FC or FCoE (Fibre Channel over Ethernet) for storage I/O connectivity.



A general industry trend is has been increased deployment of networked storage, either NAS (NFS and CIFS file sharing) or block SAN (iSCSI, Fibre Channel or FCoE). However there is also a growing market for external shared DAS and switched SAS storage solutions.

**Shared and switched SAS storage today**

While iSCSI and FC are well known as SAN server to storage I/O connectivity options (along with growing awareness about FCoE), there is another high performance option and that is SAS. SAS, also based on the SCSI command set, continues to evolve with support for faster 6Gbs speed, 10+ meters cable lengths (up to 25 meters using active cables), and a very good price to performance ratio. The combination of price, performance, shared connectivity and distances is well suited for clustered and high density blade server environments.

Common storage solution requirements:

- Maintain control of information assets resources
- Importance of data security (logical and physical)
- Performance (throughput, latency, and IOPS)
- Availability, reliability and uptime
- Capacity to support growing data needs
- Flexibility to adapt to different environments
- Cost effective and energy (Green IT) efficient
- Easy to acquire, install and use
- Storage moving closer to servers even with networked and cloud scenarios (Figure 2)

For high density scale up and scale out environments, storage is moving closer to servers in the form of shared and switched SAS. While the cloud service and the servers providing the functionality may be located some distance from servers accessing them, a good cost effective backend storage solution being used increasingly is shared or switched SAS. Another example of storage getting closer to servers is virtualization which leverages industry standard processes with external storage. Last, however not least, is consolidation, which results in servers and storage coming closer together.

**SCSI then and now**

A commonly utilized and time proven storage I/O protocol for both internal and external dedicated or shared block based storage is the SCSI command set. Prior to SANs becoming commonplace the mention of SCSI prompted thoughts of bulky parallel cabling for internal dedicated or external shared storage.

Parallel SCSI was challenged by limited storage sharing, connectivity and cabling complexity, distance limitations along with performance. Today the time-proven SCSI command set is used with several serial storage I/O interfaces including 6Gbs SAS, 8Gbs FC, IB and 10GbE (FCoE and iSCSI) for enabling flexible, scalable high performance data infrastructures.



**Making the Switch to Shared SAS Storage**

The decision about what type of server and storage I/O interface and topology is often based on cost, familiarity with available technologies, their capabilities and, in some cases, personal or organizational preferences.

In the past there was a gap in terms of connectivity or number of servers that could be attached to a typical shared SAS or DAS storage system. This has changed with the increase of native 6Gbps ports and using the SAS switch to increase the fan-out (from storage to server) or fan-in (servers to storage) number of attached servers.

What this means is that in the past to support connectivity of multiple high performance servers, 10GbE iSCSI or Fibre Channel has been used. With the advent of increased numbers of native SAS ports on storage systems along with switches, there are new options for system designers, architects and storage administrators

**The other SAS**

While this industry trends and perspective white paper is focused on shared and switched SAS for server to storage system connectivity, it is worth mentioning the other role of SAS. SAS is similar to Fibre Channel in that it is used for attaching servers to storage systems as well as being an interface option for internal and external HDD.

The importance of mentioning the other SAS role is to show the flexibility, scalability and reliability of the technology as an I/O interface being deployed in diverse scenarios including as a native or built-in capability on many servers. The business benefit of having SAS built in or pre-integrated on servers helps to reduce the cost and complexity for server to storage connectivity.

Attribute	1GbE iSCSI	6Gbps SAS	8Gbps FC	10GbE iSCSI/FCoE
Point-to-Point	Yes	Yes	Yes	Yes
Switched	Yes	Yes	Yes	Yes
Cost	Low	Low	Higher	Higher
Performance	Good	Very good	Very good	Very good
Distance	Data center and wide area	Up to 25 meters	Data center or campus	Data center or campus
Strength	Cost, simplicity, distance	Cost, performance, simplicity	Performance, scalability, distance	Performance, scalability, distance
Limitation	Performance	Distance	Cost, complexity	Cost, FCoE emerging
Servers	10's to 100's	10's	100's to 1,000's	100's to 1,000's

Table 1 – Comparing and positioning server and storage I/O SAN technologies

Table 1 compares and positions different SAN or storage sharing approaches to determine what the applicable technique, technology or tool to use for a given task. Each of the different SAN connectivity approaches can be used to do many different things, however doing so can also extend beyond its design or economic and QoS (Quality of Service) comfort zone.

### Shared and switched SAS storage scenarios

The combination of scalable modular 6Gbps SAS storage systems (for example the LSI™ Engenio® 2600 storage system), 6Gbps SAS switches, such as the LSI 6160 SAS Switch, along with ANSI T10 zoning (for optional JBOD scenarios) enables performance with low cost, safe secure storage sharing across multiple servers independent of host operating systems or virtualization hypervisors. The benefits of 6Gbps SAS are low cost, connectivity and ease of use associated with iSCSI along with the performance similar to Fibre Channel. This is what makes shared and switched 6Gbps SAS flexible enough to meet the needs of smaller environments as primary storage or as a storage device for high density servers found in enterprise, cloud or managed service providers as well as scale out storage solutions.

#### LSI Engenio 2600 Storage System



- Scale – From SMB to Cloud to HPC
- Interfaces – 6Gbps SAS, iSCSI/SAS, FC/SAS
- Scalability - Up to 192 SAS HDDs
- Flexibility - Supports 2.5” and 3.5” disks
- Capacity - Mix of SAS HDDs, self-encrypting drives (SEDs) and SSDs
- Availability – Active/active dual controllers
- Performance - 200,000 cached IOPs, 4,200 Mbs throughput from disk
- Management - Robust easy to use
- Optional Turbo Performance feature
- Security - SafeStore® encryption services

#### LSI 6160 16 port 6Gbps SAS Switch



- Scale – Transition from DAS to SAN
- Ports - 16 Mini-SAS, (4 PHY per port)
- Scalability - 1,000 SAS/SATA devices
- Flexibility - Server independent
- Form factor - Half wide pairs can be on same shelf for High Availability (HA)
- Interoperability - Server independent
- Performance - 384Gbps aggregate (16 ports each 4 PHY x 6Gbps = 24Gbps per port)
- Management - Out of band
- SAN security – ANSI T10 zoning for JBOD

In general, shared and switched 6Gbps SAS is well suited for:

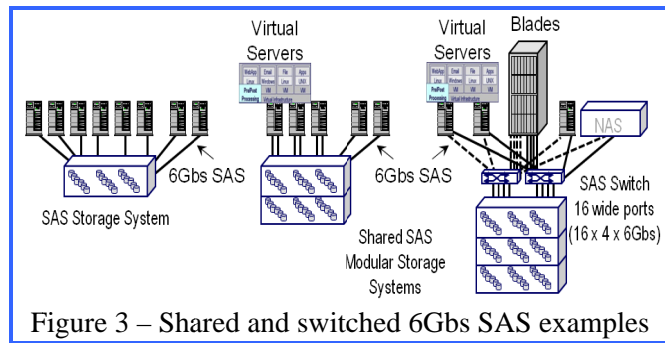
- Enabling enhanced resiliency, flexibility and scalability over traditional DAS storage
- Cost effective SAN scalability for one room, consolidated and co-located servers
- Flexibility and resilient storage options for scale up and scale out cloud services
- ROBO, workgroup, departmental and SMB affordable storage sharing

Applications or environments well suited for shared and switched 6Gbps SAS include:

- Storage for high density servers and blade systems
- Disk to disk (D2D) backup/restore and data protection appliances
- Video and multi-media streaming, security or gaming surveillance, seismic analysis
- Database, data warehouse and business analytics
- Scale out NAS, object and cloud storage solutions
- Application and server clustering (MSCS and Oracle RAC)
- Email, messaging and collaboration (Microsoft Exchange and SharePoint)
- Server and desktop virtualization (Citrix/Xen, Microsoft HyperV and VMware vSphere)

**Making the Switch to Shared SAS Storage**

Figure 3 shows configurations of shared and switched 6Gbs SAS storage supporting different application or environment needs. On the left side of Figure 3 are six single attached servers with a dual attached NAS gateway or storage appliance. A growing industry trend is the use of NAS file serving appliances that can attach to shared or switched storage, such as SAS for hosting unstructured data including virtual servers. The six servers are configured for non-high availability (HA) while the NAS storage is configured for HA.



In the middle of Figure 3 are three servers dual attached in a highly available (HA) configuration with a NAS gateway. The NAS device shown in Figure 3 leverages the shared direct attached or switched SAS storage,

supporting traditional along with virtual servers. Two of the servers, for example physical machines (PMs) in the middle of Figure 3 are configured as a cluster hosting virtual machines (VMs) for hypervisor such as Citrix/Xen, Microsoft HyperV or VMware vSphere. On the right of Figure 3, the configuration is enhanced by adding a pair of 6Gbs 16 port SAS switches along with the addition of a high density blade system. Server blades in the blade system can be configured for a mix of different applications that access the shared SAS storage system via SAS wide ports (4 PHY x 6Gbs) for performance and availability provided via HA configuration. Various applications can be deployed on the blade system, for example Microsoft Exchange, SQLserver along with SharePoint or a single scale out application requiring multiple servers and high performance storage.

**Summary**

The availability of 6Gbs shared and switched SAS gives system designers, architects and IT administrators an option for significantly boosting performance over 1Gbe based iSCSI without the complexity or cost of more expensive 8Gbs Fibre Channel, IB or 10Gbe iSCSI. In addition to the current 6Gbs, there is a SAS roadmap for next generation 12Gbs speed with backward compatibly to protect investment in current 6Gbs and past 3Gbs solutions. These and other benefits discussed in this industry trends and perspectives white paper make 6Gbs shared and switched SAS connectivity a viable option for both entry level, SMB, ROBO, workgroup or departmental as well as for high density, cloud, managed service provider (MSP), high performance and consolidation scenarios. Learn more at LSI SAS landing page located at [www.lsi.com/sas](http://www.lsi.com/sas).

**About the author**

Greg Schulz is founder of Server and StorageIO, an IT industry advisory consultancy firm, and author of the books *The Green and Virtual Data Center* (CRC), *Resilient Storage Network* (Elsevier) and coming summer 2011 *Cloud and Virtual Data Storage Networking* (CRC). Learn more at [www.storageio.com](http://www.storageio.com), [www.storageioblog.com](http://www.storageioblog.com) or on twitter @storageio.

All trademarks are the property of their respective companies and owners. The Server and StorageIO (StorageIO) Group makes no expressed or implied warranties in this document relating to the use or operation of the products and techniques described herein. StorageIO in no event shall be liable for any indirect, consequential, special, incidental or other damages arising out of or associated with any aspect of this document, its use, reliance upon the information, recommendations, or inadvertent errors contained herein. Information, opinions and recommendations made by StorageIO are based upon public information believed to be accurate, reliable, and subject to change. This industry trends and perspective white paper is compliments of LSI [www.lsi.com/sas](http://www.lsi.com/sas)