

White Paper

How To Cost-Effectively Optimize Your Investment with Serial-Attached SCSI

- SAS Advantages
- SAS Expander
- Comparison
- Case Study



How To Cost-Effectively Optimize Your Investment with Serial-Attached SCSI

Since the emergence of Serial-Attached SCSI two years ago, many high-end storage solution providers and OEMs have recognized the advantages of SAS over its parallel predecessor. Vertical markets that have since adopted SAS technology are entertainment industries, medical, financial institutions, and research & government institutions, which typically need performance and/or capacity. SAS can offer good value, performance, and scalability.

SAS Advantages

In general, everyone recognizes the following advantages SAS can offer:

- **Superior aggregated performance over 4-lane wide port -- this is like offering 12 Gb/s SAS**
- **Dual-ported design with fail-over capabilities**
- **Support both SAS and SATA HDD in one single enclosure to avoid redundant qualification and certification efforts, as well as the ability to offer tiered storage**
- **Available *edge* and *fan-out* expander technologies allowing for easier expansion and unmatched scalability**
- **Longer interconnect distance allowed compared to SATA for low-cost expansion enclosure solutions**

Although overall market share is still relatively small comparing SATA, parallel SCSI, and even Fiber-Channel, the majority of the next generation solutions currently being designed by the OEMs are based on SAS technology. Lack of full understanding on each of these advantages is being blamed for the slower-than-desired market growth rate. Among these, the advantages of what SAS Expander technology can offer is the least understood thus misleads many people when making decisions on which storage technology to choose.

A SAS Expander is basically designed to allow the Maximization of Available Bandwidth Utilization. In other words, it maximizes the value on investment.

Since SAS does not use a parallel bus to attach multiple devices to the same data path, the only way to “share” the bandwidth would be to design a “bridge device” to connect multiple devices to the available bandwidth. This bridge device is a SAS expander.

SAS Expander

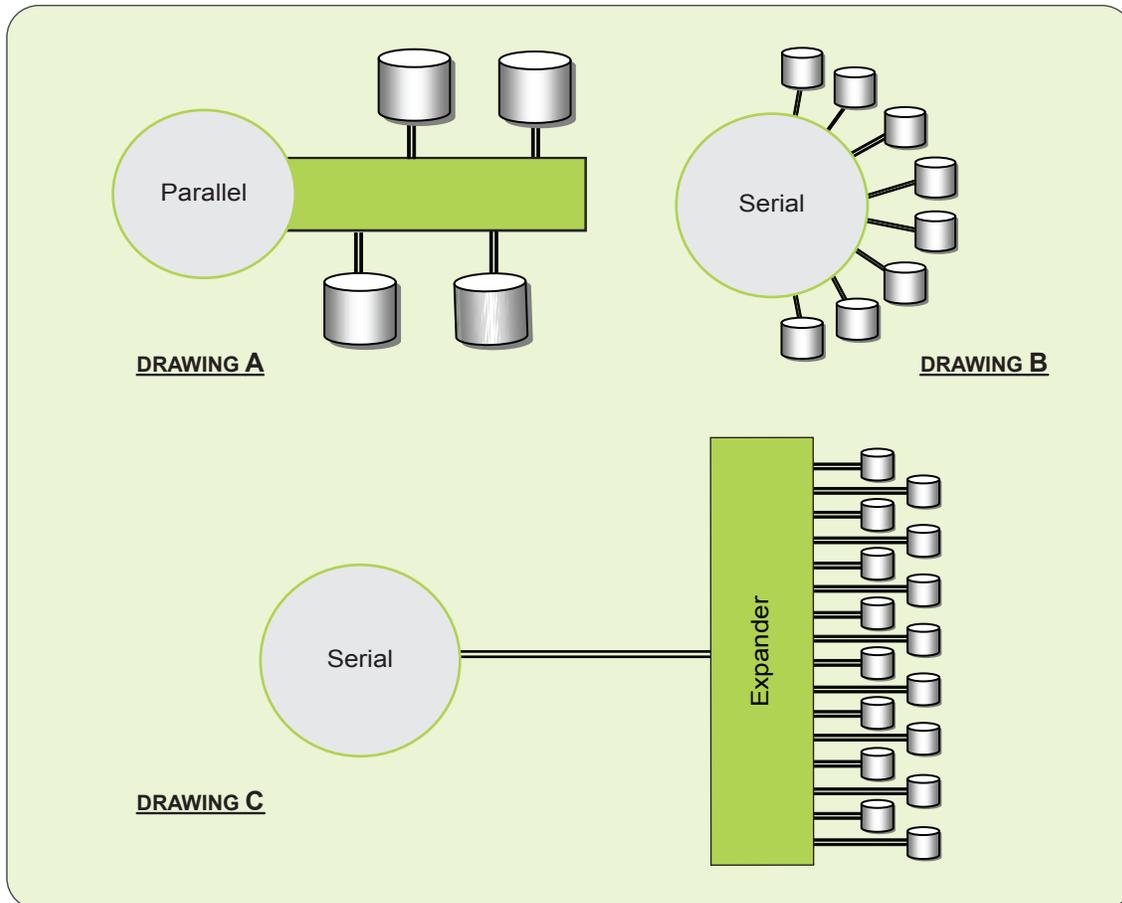


Figure 1

As we have seen, Figure 1 (Drawing B) above, a typical 8-port SATA RAID controller can offer a total bandwidth of 24 Gb/s (3.0Gb/s x8). However, one can only connect eight HDDs to this controller, even though each HDD can barely fill 20% of the available bandwidth per each port of SATA. What this means is the average throughput on a SATA 3.0G HDD is no more than 60 MB/s under the best case scenario of 100% sequential read. This is only 20% of 300MB/s each SATA port is designed to support. It is not a good way to connect more HDDs to each SATA port unless a SATA II port multiplier is used. Since SATA II port multipliers are not widely supported by many manufacturers, 80% of the total available bandwidth is wasted. This is not a good value in anyone's opinion.

In the case where large numbers of low-cost SATA disk drives are desired for applications such as Disk-to-Disk Backup or Virtual Disk Library, say 24, 36, or even 48 drives, one would either have to use a large port-count SATA controller (such as a 24-port one) or multiple smaller port-count ones.

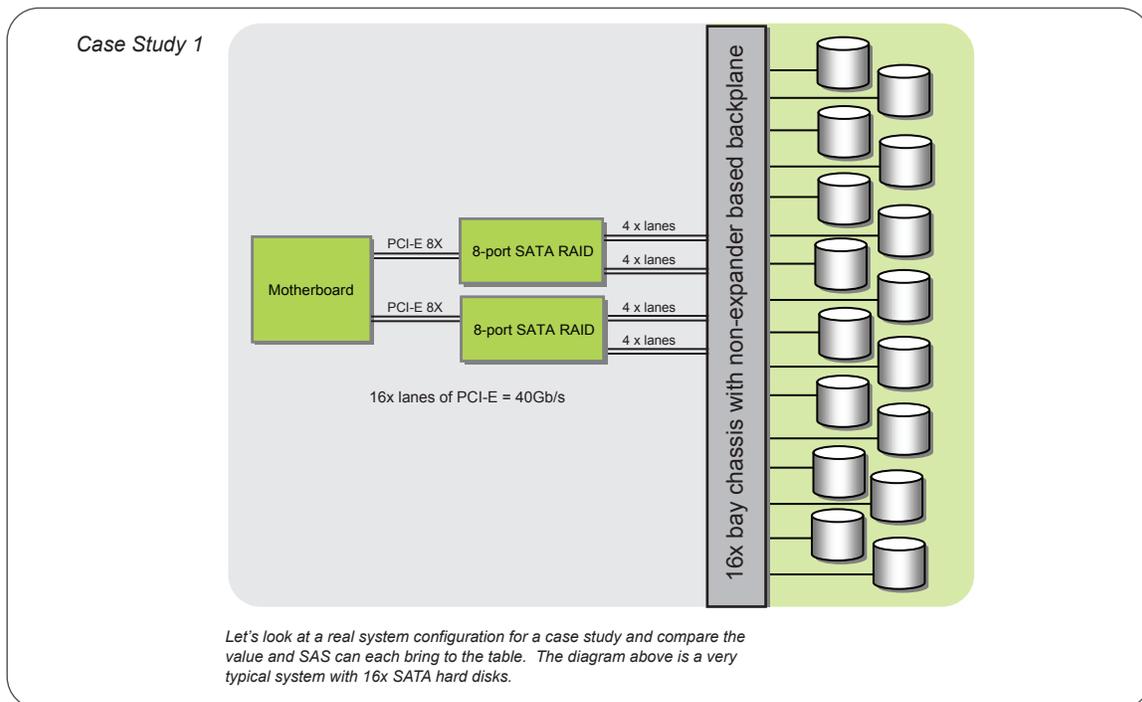


Figure 2

Cost for the total solution can add up quickly as more controllers are needed to support the drive capacity the system is calling for. Also, when using multiple controllers that do not support RAID-spanning, a software striping will need to be implemented across multiple RAID arrays to bring all the disks under one volume. This will add more loading on the server board CPU thus hampering the overall performance.

Case Study 1:

An 8-port SATA RAID controller is retailed at around \$600. A typical 16-bay SATA or non-expander based SAS storage server enclosure is priced at around \$1,500. To support 16x SATA HDD, we will need two 8-port controllers for a total cost of \$2,700 for the solution minus the drives and motherboard. Since SAS controllers now mostly come with 8x ports only, we will stay with 8-port SATA controllers only for comparison purposes. Please note two PCI-E 8x lanes for total of 40Gb/s bandwidth are occupied on the server board to support these 16x hard disks. This means each HDD is allocated around 250MB/s of bandwidth to use, even though actual throughput on the HDD is not even close to that.

Now let's look at a SAS expander based solution below:

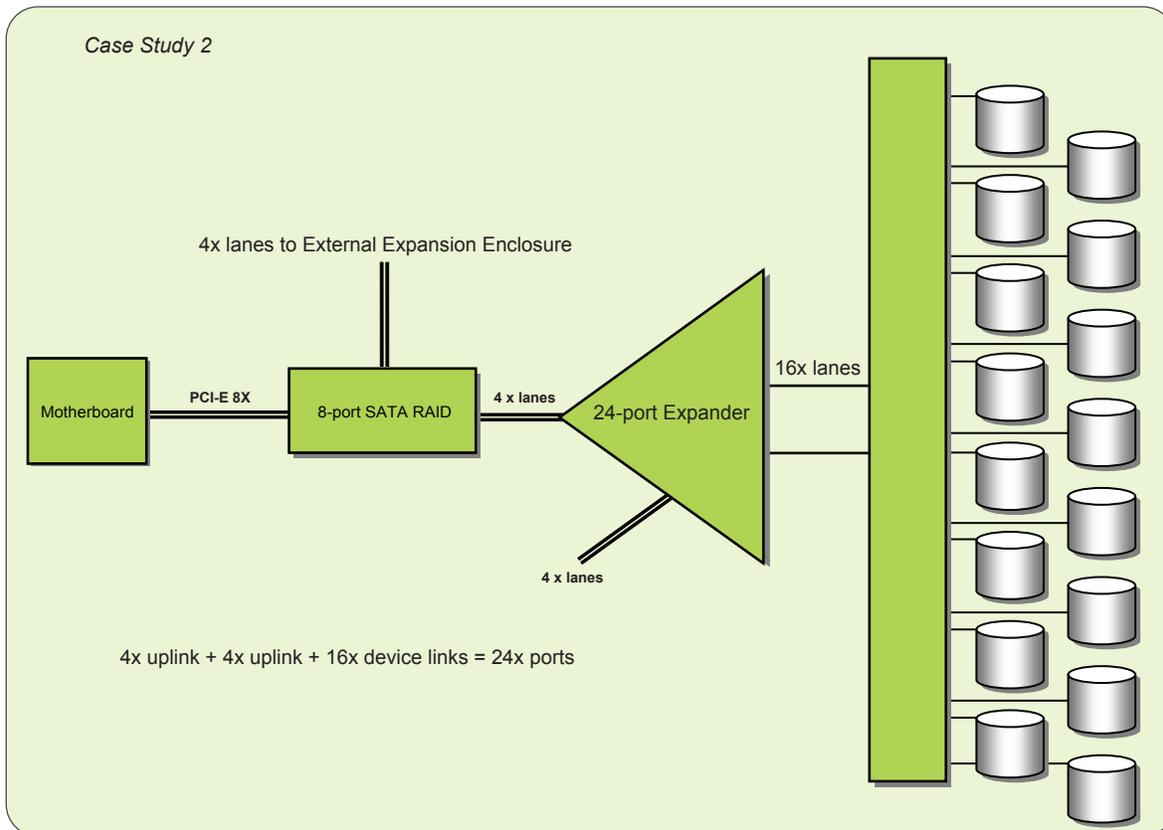


Figure 3

Case Study 2:

By using a 24-port SAS expander we only need 4x lanes (one cable) from an 8-port RAID controller, and leave the other 4x lanes for future expansion using external enclosures. Although this controller is using a PCI-E 8x slot, it is using only half of the total bandwidth of 20Gb/s to support these 16x hard disks. This is equivalent to around 1GB/s of data path (60 + MB/s per HDD allocated) which is a much better utilization of available bandwidth. Should the actual performance exceed 1 GB/s aggregating the performance of 16x hard disks, a second SAS RAID controller with 8x internal ports can be used instead to double the bandwidth, provided that the controller supports link aggregation.

Comparison

For the overall cost comparison, a SAS expander based server chassis with a single 24-port expander is retailed at around \$2,200. Together with an 8-port SAS RAID controller priced at around \$1,000, we are looking at a total cost of \$3,200 without MB and HDD. Although this solution is \$500 more than using a pure SATA solution, it leaves another precious PCI-E 8x slot open on the MB for other add-on card use. When comparing to a SAS/SATA solution using two SAS RAID controllers and a non-expander based enclosure, the net cost of this solution is actually less expensive (\$3,500 versus \$3,200). More importantly, the ability to further expand out of the same 8-port controller externally and the better utilization of total bandwidth brings much better value to solution.

Top View of Storage Server

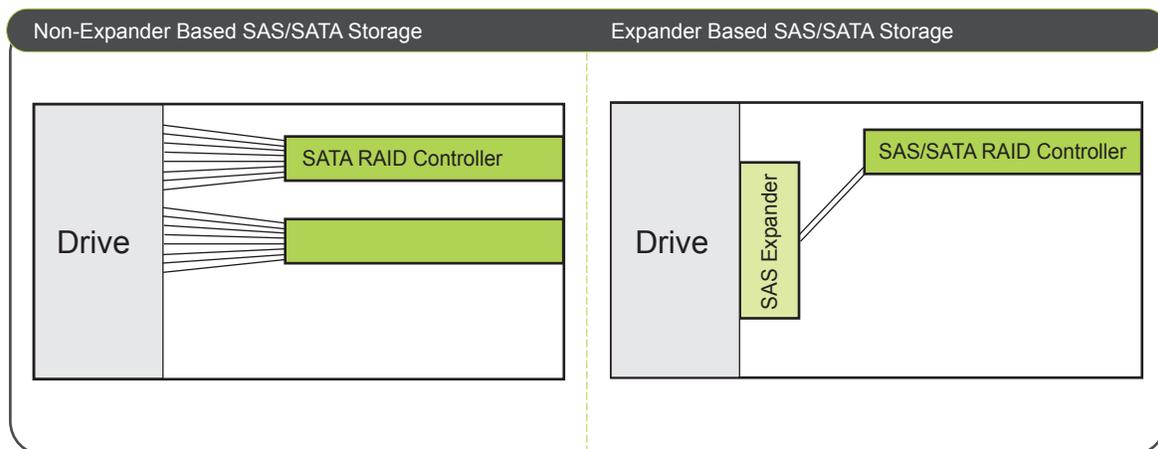


Figure 4

One last important advantage of a SAS expander is the ability to provide external expansion solutions using direct SAS links. Until this was available, users had to use FC-SATA or SCSI-SATA solutions which are all very expensive, and in most cases they are not daisy-chainable. For disk-based archive solutions where throughput is not a concern, the ability to daisy-chain from enclosure to enclosure using one single SAS RAID controllers, adds even more value to SAS. The ease of manageability across the enclosures is also unmatched by other solutions. However, there is a typical limitation of around 128 devices each controller can support, but in terms of 16x bay enclosures that is still eight enclosures a single controller can support.

The ability to recognize and take advantage of SAS technology combined with available full-featured SAS enclosures, a storage system architect can minimize the Total Cost of Ownership and maximize the value on investment.

Cost Analysis and Comparison



	Case Study 1 (Figure 2) SATA only or Non-Expander based SAS/SATA System	Case Study 2 (Figure 3) Expander-based SAS/SATA System
Performance	Two PCI-E 8X slots are occupied. (Only 1/4 of 40Gb/s bandwidth is used, 3/4 bandwidth is wasted.)	One PCI-E 8X bandwidth (20Gb/s) splits into half (10Gb/s) for internal storage, allowing 60+ MB/s per HDD bandwidth. Another half of PCI-E 8X bandwidth (10Gb/s) could be used for storage expansion to external JBOD. 2nd PCI-E 8X slot can be used for other add-on controller card.
Cost (*)	\$1,500 - SAS/SATA Non-Expander based Enclosure +1,200 - 2x SATA RAID Controllers ----- \$2,700 OR \$1,500 - SATA Enclosure +2,000 - 2x SAS RAID Controllers (MSRP) ----- \$3,700 <small>*SATA RAID Controller = \$600 *SAS RAID Controller = \$1,000</small>	\$2,200 - SAS/SATA Expander based Enclosure +1,000 - 1x SAS RAID Controller ----- \$3,200 <small>(MSRP) *SAS RAID Controller = \$1,000</small>
Others		<ul style="list-style-type: none"> • Simpler cabling enables better enclosure cooling • Maximization of available bandwidth utilization • Bandwidth-sharing using "Bridge Device" aka SAS Expander

Disclaimer: *Cost estimation is for reference only. Product specifications and pricing are subject to change without notice.

