1.0 Introduction

It used to be straightforward to decide what type of hard drive to integrate into a system. You only had the following two choices:

- ATA drives for desktops
- SCSI and Fibre Channel for enterprise systems

Hard drive interface changes and new system topologies have changed hardware choices. SCSI drives have transitioned to Serial Attached SCSI (SAS) drives, and ATA drives have transitioned to Serial ATA (SATA) drives. Desktops now support SATA drives, but enterprise class systems are based on the SAS technology that can utilize both SAS and SATA drives.

Drive vendors are now building 10K RPM SATA drives geared toward enterprise applications, as well as SATA drives geared toward desktops. To get the most out of your storage, we need to understand the differences among enterprise 15K RPM SAS drives, 10K RPM enterprise SATA drives, and 7.2K desktop SATA drives.

2.0 Introduction to Serial Attached SCSI

The traditional parallel 8 or 16-bit SCSI interfaces have evolved to the newer Serial Attached SCSI (SAS) interface. This represents the logical evolution of more than 25 years of technology development and infrastructure investment. In SAS, the SCSI protocol is transported over a serial interface. Compared to the parallel SCSI interface, SAS enables:

- Faster device interconnect speeds
- Simpler cabling
- Improved system reliability
- Preservation of existing SCSI capabilities

SAS also improves connectivity to larger numbers of drives via an Expander. An Expander is an optional portion of SAS and it enables one or more SAS host controller ports to connect to up to 64 drives.

The SAS interface standard has been adopted by the American National Standards Institute (ANSI). The Serial ATA specification extends the ATA technology by enabling disk interconnect speeds of 1.5 Gb/s (150 MB/s) and 3.0 Gb/s (300 MB/s). A wide variety of SAS and SATA products have been shipping for several years including:

- Hard drives
- Controllers
- Servers
- Networked storage systems

3.0 Introduction to Serial ATA

Serial ATA is the logical evolution of parallel ATA, the prevalent disk interface technology in:

- Desktop PCs
- Entry-level servers
- Entry-level networked storage systems

In addition, Serial ATA enables:

- Hot-pluggable devices for easier device replacement
- Thinner cables for improved in the box cooling
- Longer cables for simpler cable routing
- Cyclical redundancy checking (CRC) for enhanced data reliability

The point-to-point cabling of Serial ATA also eliminates the confusing master/slave configuration of parallel ATA.

4.0 Serial ATA -- Serial Attached SCSI Compatibility Overview

Compatibility is defined as the ability of a SAS system to use either SAS disks or SATA disks. A SATA system does not allow a SAS disk to be used.

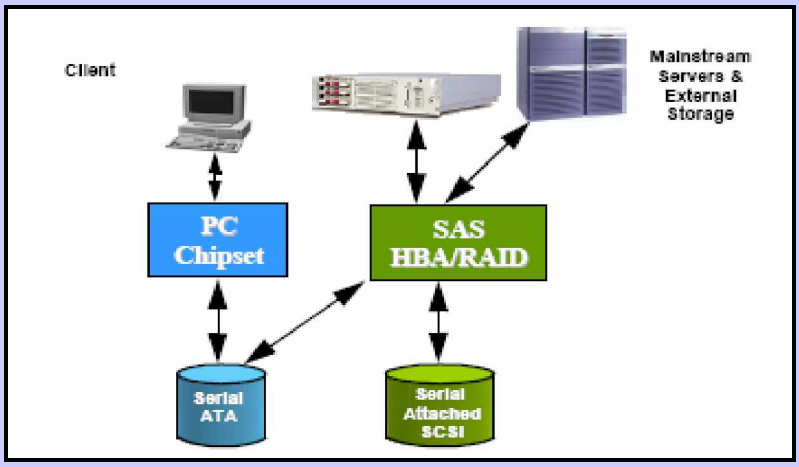

Figure 1. Serial ATA/Serial Attached SCSI Compatibility

Figure 1 illustrates how mainstream servers and external storage systems take advantage of this compatibility. A SAS drive plugs into a SAS-enabled system. The same SAS system also accepts a SATA drive via an expander or via a SAS protocol chip that also supports the ability to connect directly to a SATA drive.

Compatibility is enabled because the SAS and SATA drive connectors are very similar. This is seen by looking at the connectors. Figure 2 shows a picture of both a SATA hard drive connector and a SAS hard drive connector.

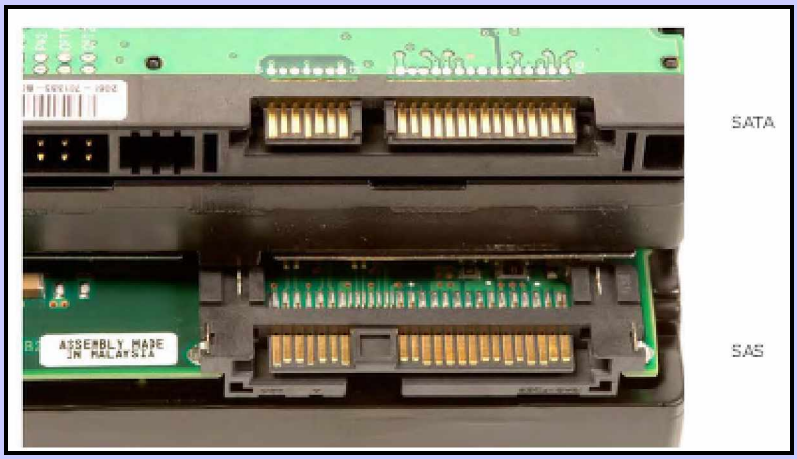

Figure 2. SATA and SAS Hard Drive Connectors

Both connectors look the same and both have the same number of pins, and the pins look to be the same size and shape. However, the SATA drive connector has a notch, while the SAS connector does not.

A system that only supports SATA hard drives uses a SATA cable or a SATA backplane connector that only allows a hard drive with a notch to be installed. Thus, this notch prevents a SAS drive from being plugged into a SATA only system. However, a SAS system uses a SAS connector that allows a hard drive with the notch (SATA) or a hard drive without the notch (SAS) to be installed. Another difference between the connectors is that the SAS connector has a second set of pins on the underside. Since SAS drives are dual-ported, these additional pins are for the second port. SATA drives, which are single-ported, do not have these pins.

5.0 Enterprise SATA drives

Historically, drives built for a desktop are built to conform to different performance and reliability requirements than drives built for a server or external storage. The specifications for SAS and SATA do not document reliability nor performance requirements. These differences are due to market dynamics.

The Enterprise SATA drive was developed to bridge the gap between a desktop SATA drive and Enterprise server SAS drives. Enterprise SATA drives have better reliability than desktop SATA drives. Despite no reliability standard, the Enterprise SATA is being defined by market needs for higher performance and better reliability at a price desktop SATA and Enterprise SAS. Enterprise SATA drives are generally more expensive the desktop SATA drives, although not as expensive of SAS drives.

For example, in November, 2007, $400 would purchase a 1TB desktop 7.2K RPM, 3 MB/sec SATA, or a 10K, 3GB/sec 146G Enterprise SATA, or approximately $500 for a 146G, 3GB/sec, 15K SAS drive.

6.0 Comparing Hard Drives

Table 1 compares Desktop SATA, Enterprise SATA and SAS hard drives.

Table 1. Comparing Hard Drive Types

Performance			
Transfer Direction	Half-duplex	Half-duplex	Full-duplex with Link Aggregation
Transfer Rate	1.5 & 3.0 Gb/sec (6 GB /sec planned)	1.5 & 3.0 Gb/sec (6 GB /sec planned)	3.0 Gb/sec (6 GB /sec planned)
Spindle Speed	7.2K RPM	7.2K to 10K RPM	10K to 15K RPM
Availability			
Number of Ports	Single-port HDD's	Single-port HDD's	Dual-port HDD's
Number of Hosts (Initiators)	Single-host	Single-host	Multi-initiator
Capacity	Up to 1TB	Up to 1 TB	Up to 300 GB
Reliability	Good	Better	Best
300GB Cost 11/07	< $100 7.2K RPM 1.5 GB/sec xfer	$750 10K RPM 3 GB/sec xfer	$900 15K 3 GB/sec xfer

As shown in Table 1, the differences between the drives break down into five categories, which will be discussed on the following pages.

- Performance
- Capacity
- Reliability
- Cost

6.1 SAS and SATA Performance

SAS drives use full-duplex techniques to information is transferred to the drive and from the drive at the same time. With a 3Gb/s SAS hard drive, the resulting link to that drive runs at 6Gb/s. A SATA drive, whether desktop or enterprise, utilizes half-duplex communication. When information is being sent to the drive, any information that needs to be sent from the drive must wait for the first communication to complete. As a result, the link to a 3Gb/s SATA drive runs at a maximum of 3Gb/s.

Hard drives are made up of platters that spin at high speeds. When accessing data on the hard drive, these platters must move to the proper position. The faster the platters move to the proper position for each read or write, the faster the action is completed. As Table 1 shows, SAS drives 15K spin rate is faster than Enterprise SATA's 10K rate or desktop SATA 7.2 RPM drives. Spin rate is so important to throughput because it always takes one complete spin to read a track.

To summarize, SAS drives have higher performance than desktop SATA drives. Enterprise SATA drives may perform better than desktop SATA drives, but because the SATA technology is inherently half-duplex, any SAS drive outperforms any SATA drive, even an Enterprise SATA drive.

6.2 Redundancy and Fault Tolerance Single Ported SATA vs. Dual Ported SAS Drives

Another key difference between all SATA drives are single ported, while all SAS drives are dual-ported. Dual porting allows a redundant system to be built. When one of the HBAs or SAS controllers fails, the other is still able to access the SAS hard drives and the data stored on those drives.

6.3 Capacity

SATA drives are offered at significantly higher capacity than SAS drives. SATA drives spinning at 7,200 RPM are built on 95mm disks while 10,000 RPM and 15,000 RPM SAS drives must be built on smaller disks. Today, the largest capacity available for SAS drives are 300 gigabytes, which means a 4-drive RAID 5 array of SAS drives offers 900 GB of storage space. That same physical space allows installation of four 1 TB SATA drives, resulting in a 4-drive RAID 5 array that offers 3TB of storage space.

6.4 Reliability

Hard drives are made of moving parts. Circular platters inside the drive rotate at high speeds, 7,200 RPM or 10,000 RPM for SATA to 15,000 RPM for enterprise SAS drives. The arm moves the heads to the proper location on the platters to read or write data. The mechanical mechanism to perform a read or a write is very precise with very small tolerances. A hard drive typically contains multiple platters that all spin at once. This rotation causes the drive to vibrate. Too much vibration impacts performance or damages the platter, resulting in data loss.

Performance is impacted when the vibration causes a head to miss the exact spot required to read or write, requiring the platter to rotate again to place the head over the data. This additional rotation takes approximately 8.3ms on a 7,200 RPM drive. The average latency of a 7,200 RPM SATA drive is less than 4.5ms, but still needing an additional rotation to find the data is a significant performance impact. Too much vibration can cause the head to hit the platter, damaging the platter and resulting in data loss. Typically, an individual hard drive is designed to handle its own vibration.

Challenges arise in systems with multiple hard drives in close proximity, such as in a storage array or a server. All the drives are spinning at once, but not spinning in synch with each other. The vibration from one drive can impact the other drives. SAS drives are typically used in multi-drive servers or storage arrays, and are designed to accommodate this additional vibration. Desktop SATA drives are designed to be in systems with one or two drives, and thus, usually not designed to handle such extra vibration.

When desktop SATA drives are used in SAS servers or storage arrays, rotational vibration may become a problem. For this reason, many server and storage array vendors do not specify desktop SATA drives for use in these systems. Thus, a system that supports both SAS and SATA drives often specifies and validates enterprise class SATA drives and not desktop SATA drives.

6.5 Price

Traditionally, SCSI drives were significantly more expensive than ATA drives at the same capacity. The transition to SAS and SATA retains similar pricing differences. It is not uncommon to see SAS drives priced more 5 times higher than desktop SATA drives on a $/G basis. While Enterprise SATA drives cost more than desktop SATA drives, these drives are less expensive than SAS drives.

7.0 Which Drive Type to Choose

To make a correct, choice of Desktop SATA, Enterprise SATA, or Enterprise SAS drives, read the system documentation for small details hidden in technical specifications pages. Overlooking something like a 1.5GB transfer rate could cause you to purchase drives that will have less performance. Some system vendors put strict guidelines in place for these products including which drives are validated and which drives to combine in a single system.

After working through the guidelines from the system vendor, choose the drive that is best designed for your intended use. In general, desktop SATA drives are best suited to client and desktop systems. The reliability of desktop SATA drives makes them ill-suited for enterprise servers or external storage systems.

Choosing between enterprise SATA and SAS drives for an important server or an external storage system is more difficult. It is tempting to focus on cost, but extra reliability, dual porting, and faster spin rates make it clear that if you go with Enterprise SATA, you will be saving money, but losing capability as a result of not choosing SAS drives.

If you need lots of storage, but only have a small budget, you may have to choose SATA. But certainly in this case, you should use a highly fault-tolerance approach, such as RAID-6, to gain back some of the reliability.

The following categories can help you make a good storage decision:

- When performance is a key attribute of the application, seriously consider SAS drives. As described above, SAS drives rotate faster than SATA drives and with full-duplex capability double the bandwidth to the drive.

- When the application requires the drive to be highly available and communicate with two HBAs or controllers at the same time, the use of SAS drives is desirable.
- When the application requires the highest levels of reliability, use SAS drives. While Enterprise SATA drives are significantly more reliable than Desktop SATA drives, SAS drives are designed and built to be even more reliable.
- When the application required large amounts of inexpensive storage, Enterprise SATA drives are indicated as a good solution. The sheer capacity available in a single Enterprise SATA drive enables very large storage arrays to be built using much less physical space and power than needed to build the same size array using SAS drives.

8.0 Summary

System builders today have unprecedented choices when deciding which hard drive best suits the specific application. SAS drives remain the best solution for applications requiring:

- Highest performance
- Highest reliability
- Most available systems

Carefully reviewing the attributes of the different hard drives available you to develop the best storage system to meet your needs.