Agilysys

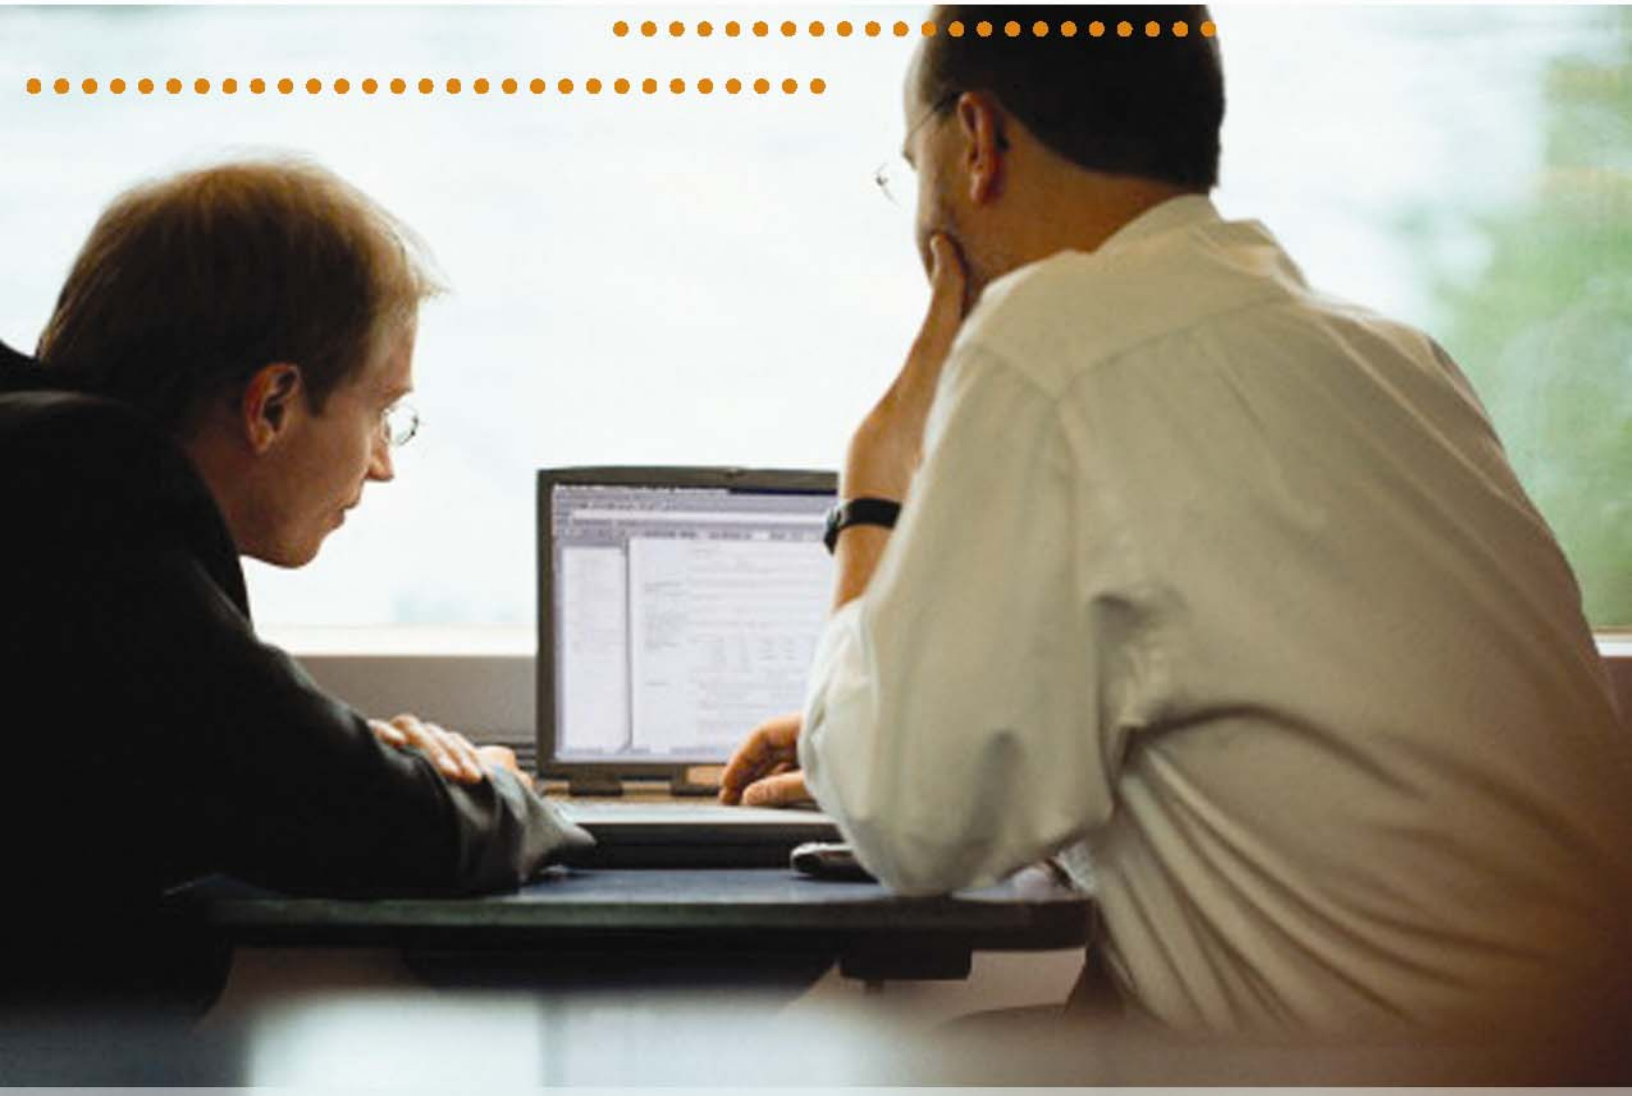

Serial ATA: Application Considerations for the Enterprise

April 2004

Serial ATA: Application Considerations for the Enterprise

April 2004

by Rob Taylor

IBM Storage Program Manager
Agilysys, Inc.

The intent of this paper is to educate the reader about information that will allow you to intelligently determine where Serial ATA fits in your environment and where it does not, and in fact where it should be avoided. It will be shown that Serial ATA solutions will not be better in performance than FC drives, but in many cases good enough to justify the lower dollar costs.

As we all know, the IT industry has been under an extreme budget squeeze for the last 30 months. CIOs no longer have a blank check to spend on upgrading or implementing new systems. The trend to do more with less, applies not only to staff augmentation, but to hardware acquisition as well. Often it is a desirable tradeoff to obtain a lower cost solution, which may not feature the highest performance, but is perfectly sufficient for the application at hand.

Within this context, companies today face storage capacity requirements that are increasing at a tremendous pace. Enterprises are generating massive amounts of data—e-mail, e-mail attachments, video, audio and databases—and IT managers are anxious to identify and purchase the most cost-effective, easy-to-manage storage solutions available.

Information, now more than ever, is also seen as a vital corporate asset with intrinsic value. Federal requirements for information storage and access mandated by The Health Insurance Portability and Accountability Act (HIPPA) and the Sarbanes-Oxley Act, as well as other legislation are accelerating companies into regulatory compliance to protect some of their most valuable and potentially dangerous information from improper use.

Fiber Channel versus Serial ATA defined by cost

For enterprise class external storage systems with large capacity requirements, high-end disk drives, such as Fiber Channel (FC) drives or SCSI drives, usually are the desired solution. They offer high performance, availability, reliability and functionality. Unfortunately, they also account for a majority of the total cost of the system. Disk drives can represent the majority of a storage system's overall cost. After factoring in the cost of controllers, enclosures and software, disk drives can represent anywhere from 30-80 percent of the storage system acquisition cost. As a result, for certain storage applications, companies are seeking cheaper alternatives to FC disks, but without sacrificing too much in performance, availability, and reliability.

Advanced Technology Attachment (ATA) technology has been available for several years as a low cost alternative to more expensive FC disk drives. Parallel ATA technology, also known as Integrated Drive Electronics (IDE) technology, was originally developed for the personal computer environment, and has been adapted for use in the mid-range Unix environment in recent years. Several manufacturers, notably EMC and NetApp, have provided disk subsystems utilizing Parallel ATA disk drives. While these systems do provide a low cost alternative to FC drives, they do possess some limitations in the areas of performance and reliability.

The most recent advancement in ATA technology is Serial Advanced Technology Attachment (SATA) interface technology. As its name implies, Serial ATA is based on serial signaling technology, unlike current IDE hard drives that use parallel signaling. Initially, Serial ATA was developed as a follow-on PC drive technology, primarily to overcome the speed and connectivity limitations of Parallel ATA.

However, it also introduces to the ATA environment certain enterprise-like functionality and features, such as command queuing and hot pluggability, that makes this technology appropriate for certain enterprise type applications.

As such, Serial ATA brings to market a new technology that can reduce costs. It also opens up a new class of storage, known as Near Line storage, for the mid- range environment that previously was only served by enterprise class drives. The performance and scalability of Serial ATA are helping it become a preferred low-cost storage interface for certain mid-market to enterprise storage applications, such as backup and archiving, or high bandwidth applications, such as video streaming and video editing.

Cost Advantages Propel Market Adoption

The advantages of serial ATA disk drives over enterprise class drives are primarily in lower costs, pure and simple. Why and how are Serial ATA drives less costly than enterprise class FC drives? The difference in cost is primarily a function of the manufacturing process of the drives. Enterprise class drives are intended for servers and multi-user systems requiring heavy-duty performance and reliability. Because of this, FC drives are typically manufactured to a duty cycle greater than 70 percent, and for a very high work rate. They are tested to actually run 100 percent of the time, with warranties that typically run five years and MTBF in excess of one million hours.

Serial ATA drives on the other hand are designed and built to a duty cycle of around 20-30 percent, which means that they demand less stringent manufacturing standards, and less costly component parts. Because of this Serial ATA drives are approximately 25 percent of the cost of FC drives, on a cost per MB basis.

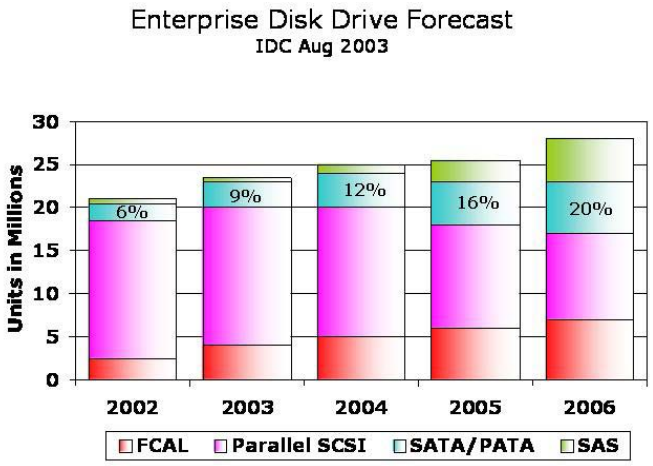

The key fact is that for some applications Serial ATA performance level is quite adequate. This fact is not lost on CIOs and other IT budget decision makers.

According to August 2003 IDC estimates* 12 percent of the drives shipped into enterprise storage solutions will be desktop class ATA in 2004. This will include both serial ATA drives and parallel ATA drives, which are still being used in large quantities. These will primarily be used for NAS storage and iSCSI. IDC also predicts this number will be growing to 16 to 20 percent over the years 2005 and 2006.

Differences between Serial ATA and its predecessor, Parallel ATA

Serial ATA is an evolution of the traditional desktop storage interconnect, Parallel ATA, or IDE as it is also known, referencing the integrated controller and disk drive technology. Parallel ATA has been the standard storage interface technology on personal computers since its introduction nearly 20 years ago. Although there have been several developments to the specification that added performance, it has changed little in the past ten years. Parallel ATA design allows only two drives per bus or two drives per cable. The classic 40-ribbon cable is large and cumbersome. These design limitations, coupled with faster applications and PC processor performance, means that it is often the cause of bottlenecks in data transfer, because it has achieved its maximum data transfer rate. At 133 MBS, Parallel ATA has reached its top end of its capability.

Serial ATA, by contrast offers many advantages that not only replace Parallel ATA, but give it many opportunities beyond the traditional desktop market as a viable option for server and NAS networked storage:

Point-to-point connectivity

Serial ATA is a point-to-point connection that allows multiple ports to be aggregated into a single controller that is typically located either on the motherboard or as an add-in, RAID card. Serial ATA features a single cable between the controller and each drive, which means that each drive can talk to the controller without having to wait for other data traffic to clear first. Unlike parallel drives that are daisy chained, each drive that is configured on a controller is individually addressable. The advantage of point to point-to-point connectivity permits faster data streaming, and easier communication with multiple disk drives.

Performance

Unlike Parallel ATA's upper limit of 133 MBS, Serial ATA technology will deliver 1.5 Gbps (150 MBS) of performance to each drive. Serial ATA defines a roadmap starting at 1.5 gigabits per second (equivalent to a data rate of 150 MBS) and migrating to 3 gigabits per second (300 MBS) by the end of the year, then to 6 gigabits per second (600 MBS). This roadmap supports up to ten years of storage evolution, based on historical trends.

Legacy OS support

Because of the legacy support inherent in the Serial ATA specification, operating system support will be simplified. The Serial ATA specification allows for additional features to be added to applications. Additional features will be subject to normal driver validation processes

Hot-pluggability

Drives on a server can be replaced without having to power down the entire machine, which saves significant maintenance time. With serial cables, they are easy to connect and disconnect.

CRC error correction

With Parallel ATA Cyclic Redundancy Checking (CRC) is performed on the data being transmitted back and forth but not on the commands. Serial ATA integrates CRC on the command and data packet level for enhanced bus reliability. Cyclic redundancy code detects all single and double-bit errors and ensures detection of 99.998 percent of all possible errors Primarily the Cyclical Redundancy is designed to be software compatible with current OS's, eliminating the need for new drives or changing existing OS.

Supports command queuing (in Serial ATA II)

Soon Serial ATA drives will be able to support command queuing, an important characteristic of FC drives that allow the drive to optimize the order of workloads.

To understand where Serial ATA best fits, it is necessary to understand the different drive classifications and characteristics:

Desktop

Parallel ATA desktop drives are commonly found in laptops and PCs, and constitute almost 90 percent of total disk drive market. They are very inexpensive to manufacture; and are designed for low duty cycle. They are not designed to take the stresses and loads of heavy utilization. Attempting to drive a laptop disk drive for 24 hrs a day, at duty cycle of 20 percent, will dramatically shorten its lifespan.

Desktop disk drives feature:

- Single-user systems, designed for lower workloads
- Lower performance and availability requirements than Enterprise class drives
- Low costs due to mass manufacturing and less expensive technology
- Designed for low activity, around 20-30 percent duty cycles
- Operating times assumed to be typically eight hours/day
- Typical warranty is one to three years

Enterprise

Enterprise class drives are utilized in high performance external storage subsystems. Enterprise drives are designed and manufactured to run 24 hours a day, and to do so without failure. These drives typically in the past have been FC or SCSI.

Enterprise disk drives feature:

- Intended for servers or multi-user systems
- High performance, availability, reliability and functionality
- Higher costs than desktop drives due to limited manufacturing (comparatively) and more robust and expensive technology
- Designed for very high activity, >70 percent duty cycle
- Typically enterprise class drives spin at 10,000 or 15, 000 rpms.
- Operating times are 24 hours/day
- Typical warranty is five years

Near Line

Near Line is a new class of drives that is taking hold in mid-range storage marketplace. Near Line drives offer similar technology from a connectivity standpoint as the desktop class. But they are manufactured and tested to much more stringent qualifications. Near Line drives are intended for multi-user systems under low workloads. As such, they offer significant improvements in drive reliability and higher capacity, though transfer rates are not as high as FC drives.

Near Line features:

- Low performance, but "similar" reliability specifications under low workloads
- High capacity, some "selected" components
- Designed for moderate activity, around 30-40 percent duty cycle
- Spin at 7,200 rpm today, with 10,000 rpm on the horizon
- Operating times are 24 hours/day
- Typical warranty is three years

Duty Cycles

One of the issues key in understanding the applicability of desktop class Serial ATA drives is realizing what the duty cycles of drives really means.

For purposes of this paper, duty cycle is the percentage of time in a 24-hour period the drive is actually doing something; searching, seeking, reading or writing data. Specifically duty cycle is the proportion of time during which a component, device, or system is operated.

The duty cycle can be expressed as a ratio or as a percentage. Suppose a disk drive actually operates for 10 seconds per minute, its duty cycle would be 10/60, or 16.67 percent. On the other hand, if an application were actually using the disk drive for 45 seconds per minute, the duty cycle would be 75 percent. In the first example, Serial ATA disk technology would probably be perfectly acceptable as a disk storage device since Serial ATA drives are designed and manufactured to a duty cycle target of around 30 percent. The second example, the duty cycle of 75 percent would probably cause excessive instances of disk failure if Serial ATA drives were used.

Application Profiles

APPLICATION	Read Intensive	Write Intensive	I/O Intensive	Throughput Intensive	Random Access	Sequential Access
OLTP	•	•	•		•	
Data Warehouse	•		•		•	
System (SCP)	•		•		•	
File Serving	•			•	•	
Medical Imaging		•		•	•	
Web/Internet	•			•	•	
Multimedia Video	•			•		•
Document Imaging		•		•		•
CAD/CAM	•		•		•	
Backup/Recovery		•		•		•

It is important to realize that the selection of a properly configured disk storage array is often driven by the type of applications to be run. The Application Storage Profile chart, which is not intended to be a comprehensive application list, characterizes the technical requirements of some large application classifications. Further discussion will help to reveal that Serial ATA is very well suited to a class of throughput applications, such as file serving, medical imaging, Internet serving, Web serving, video streaming, as well as some other non-throughput intensive applications under certain circumstances.

Comparison tests of desktop class Serial ATA drives with FC drives running under these applications, reveals that the Serial ATA drives may be as much as four to five times slower than FC drives (measured in IOPS) when using small (4KB) block sizes, and when attached to identical RAID controllers.

This would seem to indicate that desktop class Serial ATA drives are not the best option for use with applications that read and write data in small blocks, particularly if performance is a consideration. However, when the size of the data blocks that are used gets larger, this performance differential between desktop Serial ATA and FC becomes much less pronounced. At a block size of 512KB, the data transfer rates of desktop class Serial ATA and FC drives are very similar when driven by identical controllers.

What this means, for a range of throughput intensive applications where the activity consists of large block writes and occasional reads, as opposed to heavy block level activity, desktop class Serial ATA offers a sweet spot. In this type of application, typically users would need to retrieve an entire file, possibly making modifications, and writing it back. Activity is throughput intensive since there is a lot of data to read. Activity is not Input/Output (I/O) intensive since files will only be accessed infrequently.

On the other hand, desktop class Serial ATA is not considered optimal for OLTP (Online Transaction Processing) and other write-intensive applications because of the slow rotational speed of most Serial ATA drives. These would include applications, such as order entry systems or call centers that require constant, instantaneous access to data, such as databases and frequently accessed user data.

Where to use Serial ATA as Primary Storage

While Serial ATA is not normally used as primary storage in the midrange or enterprise environments, there are some applications for which SATA may well be appropriate because of the I/O demand characteristics of the application. Some of these application environments are listed below.

Bandwidth streaming applications

These include applications where transfer rate is more important that seek times, or which also feature media or rich content. For example, recording of seismic applications used to locate oil and natural gas deposits. It would be hard to justify that this activity is not mission critical to the success of the oil exploration company. Yet this application also exhibits characteristics that are write and throughput intensive, and therefore falls right into the sweet spot of Serial ATA.

Some business applications with low or limited IOPS performance

Some business application I/O performance requirements also fall within Serial ATA's 30-40 percent duty cycle comfort zone. For example, an order entry system for high-ticket items that trickle in during the day, rather than at a high rate. Here a low cost Serial ATA solution can also economically benefit from robust controllers, enclosures and management software that is shared between FC and Serial ATA drives.

Where to use Serial ATA As Near Line (secondary) storage

The performance and reliability characteristics of Serial ATA are much more suited to use as Near Line secondary storage. It can also be used to cache online storage for quicker backups to tape. Secondary storage represents a large percentage of most companies' data requirements and is ideally a much better fit for Serial ATA technology.

Disaster Recovery typical site

Disaster Recovery strategies requires users to mirror data from one location to another for back up. Storage at the second location may not necessarily have the same stringent performance and access requirements as the primary site. Here it would be much less expensive to put Serial ATA at the secondary site.

Tape cache

Serial ATA drives boast excellent sequential read and write performance, which makes it ideal for backup and recovery applications in which data is written sequentially to, and recovered sequentially from disk. This application is characterized by very low activity. Data is written once in large blocks, read once, and then backed up to tape.

Fixed content/managed retention data

Content managed data solutions almost always use Serial ATA in their applications. This is data that needs to be stored in a nonerasable and nonrewritable format and is most often stored in archives and repositories. As organizations address new regulatory and legal compliance requirements, they must ensure that necessary documents and records are available in their original formats and accessible when needed. These include temporary workspace applications that are written and occasionally "retrievable for reference". E-mail archives, financial archives, medical files, land records fall into this category.

Hierarchical storage management

Hierarchical Storage Management involves the movement of data from primary disk storage to some type of secondary storage based on the time the data has been in existence or data access requirements. For example, customer records may be kept on primary disk storage for rapid access for a period of time, after which the records would be migrated to some sort of less expensive secondary storage where the records are still accessible but cannot be accessed as rapidly as when they were stored on the primary disk This storage would be ideally less expensive than the primary storage. The records may ultimately be migrated to a tertiary storage device for archiving, or perhaps discarded after a period of time. Under most conditions, Serial ATA could be a good choice in this environment when used as the secondary storage in the hierarchy.

Conclusion

Industry experts expect that FC drives will almost certainly remain the dominant host connectivity standard for some time to come, but say that Serial ATA-based solutions will continue to carve out unique niches.

If you are considering Serial ATA, you must insure that you match applications to technologies, and fully understand the inherent workload requirements. Certainly, if users try to execute a direct replacement for high performance FC drives and plug in an ATA drive, they are not designed for that kind of usage. If the application is high use, high workload, and high reliability, FC drives are advisable. With Serial ATA in those environments, users must be prepared to accept the risk of potential data loss.

On the other hand, data storage implementations best suited to use Serial ATA technology reside at the Near Line or secondary location within the networked storage hierarchy. Here users can take full advantage of the cost savings and contribute well to corporate objectives of "accomplishing more with less." This type of carefully considered strategy will continue to fuel adoption of Serial ATA technology into the marketplace.

About the Author

Rob Taylor is a Storage Manager for Agilysys, Inc. Agilysys is one of the foremost distributors and premier resellers of leading enterprise computer technology solutions from HP, IBM and Oracle, as well as other top manufacturers. Agilysys has a proven track record of delivering complex servers, software, storage and services to resellers and corporate end-user customers across a diverse set of industries.