SAS and SATA Disk Drives in SAS Expander-Based Storage Solutions

xyratex

Notices

The information in this document is subject to change without notice.

While every effort has been made to ensure that all information in this document is accurate, Xyratex accepts no liability for any errors that may arise.

Bob Hansen

Issue 2.0 | October, 2007

xyratex

Contents

Introduction

The following paper describes the basic characteristics, capabilities and limitations of Serial Attached SCSI (SAS) and Serial ATA (SATA) disk drives as they are used within SAS Expander-based storage solutions. Drive performance and applications are discussed including performance considerations when SAS and SATA drives are present in a single SAS domain. Finally, the subject of rotational vibration interference (RVI) in a mixed SAS/SATA drive type environment is addressed.

The Basics of SAS/SATA and SAS Expanders

Unlike SCSI or Fibre Channel (FC) technology based storage solutions, RAID solutions and SAS Expander-based Bunch Of Disks (EBOD) based on SAS Expander technology can directly support both high performance SAS and/or SATA disk drives in the same enclosure.

SAS drives are based on the same disk drive mechanism as FC and SCSI drives, and exhibit similar reliability, performance and cost (at least in the near term). SATA disk drives, originally developed for the PC industry, are significantly less expensive on a $/Gbyte basis but do not achieve the same level of performance and reliability as FC, SCSI or SAS. A new class of "nearline" drives using either FC or SATA interfaces provides a compromise between these two.

In general – for external storage applications, SAS drives are best suited for transaction centric applications requiring high input/output per second (IOPS) performance and a high degree of reliability when subjected to the stress of high duty cycle operations. The classic example of this environment is a banking data base environment.

SATA drives, on the other hand, are much more cost effective than SAS when used in applications requiring very large amounts of storage with lower IOPS performance requirements and a relatively low duty cycle. Examples would include nearline file archives, virtual tape libraries, large file applications such as video storage, and anywhere large amounts of "warm" data storage is required. The latter "warm" data storage examples would include end user e-mail and image data being accumulated at an amazing rate at companies such as Yahoo and Google today.

Applications for SAS and SATA Drives in a Single SAS Domain

Clearly, some applications can achieve the best possible cost/performance point by using a combination of SAS and SATA drives. For example, in an e-mail application, the images and messages most recently created is the information most likely to be accessed. This data could be stored on SAS drives to provide the best possible performance while the vast majority of information, the data rarely accessed, could be stored on SATA drives where the cost per Gbyte stored is at a minimum. SAS Expander-based storage solutions directly enable this type of mixed SAS and SATA drive environment, allowing either drive type to be supported in the same SAS domain and in the same disk enclosure.

xyratex

With SAS, all of the disk drives in a SAS domain (single address space) whether SAS or SATA, are presented to the initiator (host) as virtual SCSI devices, making it very straight forward to perform basic operations. With relatively low performance requirements or single threaded applications, the user should have few problems working with mixed domain configurations. As the demand for performance is ramped up in applications requiring simultaneous access to both SAS and SATA drives in the domain, the application must be aware of the difference in drive performance characteristics to maintain deterministic access patterns.

Performance Considerations in Mixed SAS/SATA Drive Applications

The performance characteristics of SAS and SATA drives are very different. SAS drives, like parallel attached SCSI and Fibre Channel devices, are designed for server and workstation applications. SCSI protocol, which is also the foundation for Fibre Channel, was designed for high performance, multiple drive, multiple initiator applications. The mechanical and electrical aspects of the SAS disk drive are also designed to deliver high performance with low response times for data read/write requests.

In contrast, SATA is designed for PC applications where the lowest possible cost per GByte with acceptable performance is the design goal. SATA protocol assumes that a single drive is connected to a single host. Because of the relatively low cost disk drive mechanism employed, the very high data density on the disk and the relatively simple communications protocol, SATA drives take much longer to respond to a data request than SAS disk drives. Another important aspect of the SATA protocol is that the drive assumes that it "owns" the connection. Once a transaction is started, the drive will maintain ownership of the connection until the transaction is completed, even though it may take a long time to prepare for the actual data transfer. When a large number of drives are configured in a single domain, this feature limits overall throughput by preventing other disk drives from moving data until the transaction is completed. The recent introduction of Native Command Queuing (NCQ) capability in SATA drives has improved drive performance by as much as twenty percent in some applications.

The SAS protocol defines a mechanism that has been implemented in the SAS Expanders to guarantee fair access between drives in a domain. In practice, this works quite well when a SAS domain is made up of all SAS or all SATA drives. When both drive types co-exist in a single domain, the large differences in response time combined with the tendency of SATA drives to lock up a connection can play havoc with performance expectations.

The vendor community is aware of this problem and future generations of host bus adapters and expanders will help to minimize the problem. For the near term, applications that strive to optimize solution cost and performance by using both SAS and SATA drives in a single domain must be aware of the performance implications and optimize data access patterns accordingly.

Mixing SAS and SATA Drives in a Single Enclosure

In addition to the communications protocol, SAS and SATA drives differ considerably in their mechanical design.

In order to achieve the fast data track seek times that enable low response times, SAS drives like SCSI and FC drives, first must operate at a higher RPM than SATA (generally 10,000 or 15,000 RPM compared with 7,200 RPM for SATA), and are designed to move the drive read/write head very quickly.

In contrast, SATA drives being designed for low cost / high data density move the drive head more slowly and take much longer to align the head with the very narrow data track on the drive platter. Once on track they can deliver data almost as quickly as SAS, making SATA drives an excellent value for data streaming applications.

Within a disk enclosure, the mechanical vibrations set up by disk rotation or head seek operations in one drive can interfere with the operation of other drives. This problem manifests itself to the user as a reduction in disk throughput performance. This is generally referred to as rotational vibration interference. The worst case data access pattern for RVI is random causing maximum seek operations.

When a disk enclosure contains either all SAS drives or all SATA drives, the RVI problem is minimized due to the similar drive operating characteristics. Poorly designed enclosures can result in dramatic performance problems under certain workloads. As an example, low cost JBOD enclosures designed for direct attached workstation applications can not be expected to perform well on high performance servers. The higher more random data requests created by the server could very well set up RVI conditions that could severely limit throughput.

When a single enclosure contains both SATA and SAS disks, the RVI problem can be much worse. SAS disks create more vibration because they operate at a higher disk RPM and move their heads much more vigorously than SATA. To compensate for this, SAS drives when compared to SATA have much wider data tracks on the disk and have stronger magnets on the disk head seek mechanism. Both help SAS drives find and hold the head over the data track. The electronic “servo” capability is also improved.

SATA drives have a much more difficult problem. To reduce cost, SATA magnets are much less powerful than those found on SAS drives. To optimize capacity data tracks on the SATA drive platter are very narrow. Both of these factors make it more difficult to find a data track on the disk and hold the head over the track.

When both SATA and SAS drives are present in a single enclosure, SAS drive performance is generally not affected by the RVI created by SATA drives. In contrast, in a poorly designed enclosure, SATA drive performance can be severely degraded by the presence of SAS drives. In the worst case, with all SAS drives being continuously accessed with a random data pattern, SATA drives may not be able to function at all in this environment.

The RVI Solution

SAS and SATA drives can co-exist peacefully within the same enclosure but this does not happen by accident. The enclosure must be designed to mitigate the effects of RVI and most important, the enclosure must be tested and characterized to fully understand the capabilities of the design. The physical location of drive types may be restricted within the enclosure depending on workload and performance requirements. A detailed RVI test report is important to have for a storage enclosure being considered for any application, but it is critical for mixed SAS/SATA drive applications.

Conclusions

- SAS drives like Fibre Channel and SCSI are optimized for performance and reliability.
- SATA drives are optimized for cost and data density, and can be very cost effective in many enterprise applications today.
- SAS Expander-based storage solutions can support both SAS and SATA drives in the same address domain and in the same enclosures.
- Many applications can be optimized for cost/performance through the use of both SAS and SATA drive types.
- When using SAS and SATA drives in the same SAS domain without careful consideration by the solution architect, drive type performance differences may result in unexpected and nondeterministic system throughput.
- RVI can be an issue with any disk drive enclosure. Mitigating the effects of RVI when SAS and SATA drives are present in a single enclosure can be especially challenging.
- Proper disk enclosure design and characterization is especially critical when using SAS and SATA drives in a single enclosure.

UK HQ
Langstone Road
Havant
Hampshire PO9 1SA
United Kingdom

T +44(0)23 9249 6000
F +44(0)23 9249 2284

www.xyratex.com

ISO 14001: 2004 Cert No. EMS91560