



SAS Starter Kit

Mixing SAS and SATA Drives
in a Single Enclosure



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.

© 2007 Xyratex (the trading name of Xyratex Technology Limited). Registered Office: Langstone Road, Havant, Hampshire, PO9 1SA, England. Registered number 03134912.

No part of this document may be transmitted or copied in any form, or by any means, for any purpose, without the written permission of Xyratex.

Xyratex is a trademark of Xyratex Technology Limited. All other brand and product names are registered marks of their respective proprietors.

Issue 2.0 | September, 2007

SAS Starter Kit: Mixing SAS and SATA Drives in a Single Enclosure

Serial Attached SCSI (SAS) has begun to play a key role as the enterprise continues to move towards more optimized, cost-effective storage platforms. Overcoming many of the limitations of traditional parallel SCSI implementations, SAS solutions that offer connectivity choices and support both SAS and Serial ATA (SATA) drive technology enable enterprises to easily create tiered storage environments. However, concerns about performance expectations and Rotational Vibration Interference (RVI) in mixed SAS/SATA environments are preventing many organizations from adopting SAS technology.

SAS is the latest storage interconnect technology based on the venerable SCSI protocol; it is based on a switch based storage fabric that supports high performance, SAS and high value SATA disk drives. Initially focused on the server and direct attached storage market, SAS took the best from SCSI, Fibre Channel and iSCSI, then added additional features to produce a high speed, switch based, flexible, standards based, storage interconnect technology.

Technology Comparison

The most common disk connect technology today is PATA (Parallel ATA). Until recently PATA was found in every PC although, like SCSI, PATA is now being replaced with a high speed serial interconnect – Serial ATA. Basic SATA has very limited networking and no HA (high availability) features, this limits its usefulness in business applications. However, SATA disk drives, connected to a host through a SAS expander based infrastructure, can be a very effective implementation for today's business applications.

Below is a quick summary of key features of SCSI based storage interconnect options:

- Parallel SCSI – This low cost interconnect supports a maximum of 15 devices and is limited by the electrical performance of parallel interconnects. The large cables are difficult to manage, are limited in length and interfere with air flow within drive enclosures.
- Fibre Channel – This is the most costly option, and generally used for enterprise class systems. HA features, include dual ported, high performance drives and support for multiple hosts. Today, the FC arbitrated loop link supports up to 126 devices and operates at a rate of 4Gb/s, FC is also designed to support both copper and optical interconnect, and, when operated in fabric mode supports complex storage networking applications.
- iSCSI – Based on the IP fabric, iSCSI has an excellent networking capability and enables long distance data transfer. However, iSCSI performance is currently limited to 1 Gb/s and can not connect directly to disk drives.

- SAS – This high performance (3Gb/s with 12Gb/s defined for host interconnect), medium cost storage interconnect standard is designed for enterprise class systems consisting of multiple hosts that can scale to hundreds of drives as well as the less demanding SMB applications. SAS also supports HA features, which include dual ported disk drive connections. A key advantage of SAS connectivity is the support for both high performance SAS devices and high value SATA devices.

SAS interconnect technology with SATA drives is the best choice when both low cost and large capacity is required. Native SATA connectivity is an alternative in this case, although this solution is limited in drive count and HA capabilities. For companies who are looking for low cost and high availability, the use of HA SAS enclosure architecture ensures no single point of failure with SAS drives or SATA drives with the application of a SATA port selector device. Although FC is an alternative here, the cost point is much higher and there is added complexity because, at least for the short term, a proprietary bridging technology is required to connect SATA drives to FC. SAS connectivity is also the top choice when the application requires both high performance and high value disk drives. Again, FC is an alternative here, but at a higher cost and with non-standards based SATA bridging, SAS is a more cost effective choice.

Choosing the right drive for your application

Although there are some variations, there are two categories of disk drives being used in business applications today.

High performance drives, with native FC, SCSI or SAS connectivity are designed for high duty cycle applications and deliver the best IOPS (I/Os per second) performance. In addition, FC and SAS drives are dual ported allowing HA configurations with no single point of failure to be implemented. At any point in time the highest capacity high performance drives are significantly smaller, and more expensive than the high value drives.

High value drives, with PATA or SATA connectivity, are designed for lower duty cycle applications, have very good streaming data performance but deliver comparatively slow IOPS performance. These drives are single ported, although HA configurations can be supported with the addition of a SATA mux (port selector) device. These drives ship at very high capacity points, achieving 1TB in a single 3.5 inch drive in 2007 and are very low cost.

How do you choose?

Your choice of disk drive and disk enclosure configuration depends on application performance requirements, on-line availability and budget. Most application solutions would be best served by using both types of drives.

Logical choices include:

A high performance on-line transaction processing system such as a state lottery application, which is measured by response time and has a workload typified by many small data block transactions, clearly requires high

performance drives. In this case the much higher \$/GB of SAS or FC drives is justified given the fast response time requirement.

On the other hand, an on-line library application such as Wikipedia, which is measured by breadth of coverage and has huge storage requirements for data that is rarely accessed, is clearly an application for high value drives which offer higher capacities at a lower cost.

In these two cases, there is a clear choice of disk drive technology for the primary application. However, if we consider the complete application solution a combination of drive types is called for.

In example 1, as the lottery close date approaches there will be a tremendous transactional load on the system that can only be serviced through the use of the highest performance disk solutions. On the other hand, after the close, all of that transaction data must be preserved where it will be accessible, prior to being moved to tape. This is an ideal application for high value drives.

In example 2, the Wikipedia type application, most of the data is rarely accessed but from time to time hot spots develop requiring very fast response times to satisfy a temporary but very heavy load. In this case a caching system using relatively few high performance drives could be utilized, which would make a huge improvement to response times for the hot data set.

In the past, because of technology limitations, only the largest application installations could afford this tiered storage architecture approach. SAS based storage infrastructure enables low cost solutions to be optimized for cost and performance by using both high performance and high value drives, sometimes in the same disk enclosure.

One application that even the smallest businesses must deal with is email. This is an example of an application for which a small amount of data is accessed frequently for a limited amount of time. Data access frequency falls off quickly, but taking any of this data off-line to tape is unacceptable. This is a perfect application for mixed drive architectures. The amount of high performance SAS drive storage scales with the number of employees while the high value, SATA drive storage scales with both time and the number of employees utilizing the system.

In the smallest businesses the storage configuration could start with as few as two SAS drives and just 4 or 6 SATA drives in a single enclosure (possibly a single appliance). As the business grows SAS and SATA drives would be added as required, eventually splitting SAS and SATA into separate enclosures to make management of the data easier.

Considerations when mixing SAS and SATA drives

The adoption rate of mixing high performance and high value drives in the same enclosure has been slow. This is largely due to education about the benefits of mixing drives and the overall implementation challenges. High performance drives generally rotate at a higher Rotations per Minute (RPM) and have much bigger magnets to

support high performance seek mechanisms. High value drives that are cost focused, rotate at a slower RPM. These drives have a much higher track density and implement a cost optimized seek mechanism that is much less tolerant to mechanical vibration. If the vibration from high performance drives is transferred to the high value drives, issues can occur. Seek time can go way up on the high value drives resulting in reduced, non-deterministic performance.

RVI problems happen most often when end users mix drives in poorly engineered commodity enclosures. Properly engineered disk enclosures are designed and tested to accommodate these vibrational changes; thus, minimizing the risk of potential RVI issues.

Conclusions

SAS connectivity solutions, using SAS and/or SATA drives, provide for the first time the foundation for low cost, standards based, high performance storage solutions for most business applications. Cost optimized storage for most applications require the use of both high performance and high value disk drives. With native support for both SAS and SATA drives, SAS connectivity is the best choice for these applications. Properly designed disk enclosures can support a mix of SAS and SATA drives without performance degradation, enabling highly optimized storage solutions for small business applications.

About Xyratex

Xyratex is a leading provider of enterprise class data storage subsystems and storage process technology. The company designs and manufactures enabling technology that provides OEM and disk drive manufacturer customers with data storage products to support high-performance storage and data communication networks. Xyratex has over 20 years of experience in research and development relating to disk drives, storage systems and high-speed communication protocols.

Founded in 1994 in a management buy-out from IBM, and with its headquarters in the UK, Xyratex has an established global base with R&D and operational facilities in Europe, the United States and South East Asia.

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



Cert. No. FS29805 ISO 9001

©2007 Xyratex (The trading name of Xyratex Technology Limited). Registered in England & Wales. Company no: 03134912. Registered Office: Langstone Road, Havant, Hampshire PO9 1SA, England. The information given in this brochure is for marketing purposes and is not intended to be a specification nor to provide the basis for a warranty. The products and their details are subject to change. For a detailed specification or if you need to meet a specific requirement please contact Xyratex.