

SAS technology says “goodbye” to early growing pains

Serial-attached SCSI (SAS) is the extremely flexible successor to parallel SCSI and overcomes the performance and scalability limitations of its long serving predecessor. SAS remains compatible with the SCSI command set providing a clear migration path for existing SCSI users. The true potential of SAS resides in its flexibility; although not as robust as Fibre Channel (FC) the SAS protocol offers levels of performance and reliability that meets the demands of most enterprise users. The key differentiator over FC is the ability of SAS based systems to mix SAS drives with lower cost Serial-ATA (SATA) drives within a single storage subsystem. The prospect of Information Lifecycle Management (ILM) within a single drive enclosure becomes a reality, something not possible with the mixed FC/SATA solutions on the market today.

Most major storage vendors released mixed SAS/SATA storage solutions during 2007 but very few were able to offer complete drive intermix within a single array enclosure. Many vendors were guarded over their early struggles to mix the two technologies although it is known that for some the mechanical vibrations from 15,000 rpm SAS drives caused problems when placed alongside slower 7,200 rpm SATA drives that have a far lower vibration tolerance. In the race to release SAS/SATA storage arrays in 2007 many vendors housed their SAS and SATA drives in separate enclosures or placed restrictions on the placement of drive types within a single enclosure. It is expected that most vendors will resolve issues related to mixing SAS and SATA in a single storage enclosure by mid 2008.

SAS/SATA positioning in the data centre

The SAS interface and drive mechanical construction provides levels of performance and reliability comparable to that of FC, however SAS is not a fabric technology and is intended to complement rather than replace FC in the data centre. Conversely SATA has origins in desktop computing and maintains a price point and reliability level to compete in that highly competitive market. Both SAS and FC drives are designed for continuous 24x7 operations whereas SATA drives are designed around a less arduous 8x5 duty cycle hence their mechanical construction uses fewer stabilising parts to maintain the lower price point. The more economic construction of SATA drives also accounts for slower rotational speeds which typically lag SAS spin speeds by 50%, although SATA capacity currently runs at double that of SAS with the 1TB SATA drive

making its debut in the fall of 2007. It is also important to consider that MTBF figures for SAS and SATA drives are usually quoted with reference to their respective duty cycles making a like for like comparison a speculative calculation. Despite the lower performance and less robust construction SATA certainly has its place in the data centre assuming the limitations are understood.

Realising the benefits of the SAS/SATA – Achieving ILM in a box

Enterprise data workloads align into three broad categories and SAS certainly has the characteristics to cater for all although it is not always the most cost effective choice. Online transaction processing typical of financial applications involves a demanding volume of random small data block read/writes and will benefit from high performance SAS. SATA does not suit transaction intensive applications due to the lower rotational speed and the slower introduction of command queuing; a feature that has been available with SAS drives from the outset to optimise the number of rotations needed to serve multiple simultaneous data requests.

SATA technology may be suitable for the second broad workload category classified as throughput intensive; this is typical of media streaming applications that require high throughput of large block read/writes. The sequential nature of this workload allows SATA to be a viable contender when absolute throughput is not mandatory.

SATA delivers the most attractive price/GB for reference system workloads that operate on fixed or archival data requiring large block sequential read/writes in a non real-time environment. Typical scenarios are data mining workloads such as the replay of seismic data and disk based archive such as disk-to-disk (D2D) where disk replaces tape and disk-to-disk-tape (D2D2T) where SATA drives complement tape.

The challenge and opportunity for system administrators looking to deploy mixed SAS and SATA is the ability to seamlessly move data between different tiers of disk within a single storage enclosure; transaction intensive data residing on SAS drives, while data with more modest performance requirements resides or migrates to SATA drives. To take full advantage of the SAS/SATA mix administrators should strive towards implementing a policy based data management environment that automates data migration from SAS to SATA as data approaches specified age, capacity or performance thresholds.

Consequences of continuous capacity growth: Massive SATA drives equal lengthy rebuilds – RAID 6 to the rescue

The SAS interface will continue to evolve from its current data transfer rate of 3 GB/s, with 6 GB/s on the roadmap. Drive speeds cannot continue to increase indefinitely and will start to flatten as rotational speeds reach physical limits, manufacturers will continue to squeeze out more performance with techniques such as native command queuing as the laws of physics close in. The focus to increase areal density continues to push SATA drive capacities higher which in turn increases the probability of bad block failures which are not uncommon.

Ever-increasing SATA drive capacities result in longer re-build times when a RAID protected system recovers from bad blocks or a drive failure, and this in-turn increases the chance of a second drive failure during that time.

A typical RAID 5 rebuild can take from a couple hours to several days depending on system load and priority settings. During the rebuild the array is susceptible to a second drive failure, which in a RAID 5 configuration would result in data loss. This increased potential of hitting a second bad block during a rebuild highlights the need for RAID 6 where high capacity SATA drives are used to store business critical data. RAID 6 provides an additional level of protection through a double parity algorithm that can continue to provide data recovery in the event of a second drive failure and is now becoming a 'must-have' data protection technology in the storage marketplace.

Conclusion

Mixed SAS/SATA storage systems provide a way for companies to implement tiered storage at a more attractive price than ever before. The benefits of SAS and SATA interoperability will extend from the desktop to the data centre and provide investment protection in data management software, increased performance and greater device addressability. Customers now have the flexibility to deploy solutions that maximize the use of both SAS and SATA allowing them to adapt storage to their needs quickly and easily. The administrators challenge will be establishing the right data management policies to ensure the correct drive technologies are matched to characteristics of each data workload.

About Dot Hill Systems

Dot Hill is the market leader in providing flexible storage offerings and responsive service and support to OEMs and system integrators, from engagement through end of life. Founded in 1984, Dot Hill has more than two decades of expertise in developing high-quality, competitively priced storage products. Focused on delivering global 24 X 7 X 365 technical support, the company has more than 150,000 systems in use worldwide and numerous OEM and indirect partners. With its patented technology and award-winning SANnet II®, RIO Xtreme™, and R/Evolution™ families of storage and its Dot Hill Storage Services, Dot Hill makes storage easy.

Headquartered in Carlsbad, Calif., Dot Hill has offices in China, Germany, Israel, Japan, Netherlands, United Kingdom and the United States. More information is available at www.dothill.com.

© 2008, Dot Hill Systems Corp., All rights reserved. Dot Hill, the Dot Hill logo, the R/Evolution logo, R/Evolution, are all trademarks of Dot Hill Systems Corp. This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Dot Hill. Dot Hill reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. No part of this document may be copied or reproduced without written permission of Dot Hill Systems Corp. 7.25.08.