WHITE PAPER

HITACHI
Inspire the Next

Why SAS Matters

Cost-effective Performance with SAS Technology from Hitachi Data Systems

By Mark Adams and Alan Benway

September 2010

Table of Contents

Executive Summary

For more than 10 years, Fibre Channel disk drives have been the standard for enterprise data storage. This paradigm is changing. Serial Attached SCSI (SAS) is replacing Fibre Channel drives in high performance applications, such as online transaction processing, real time analytics and streaming media. Disk drive vendors have already signaled their intentions to ramp up development of SAS drives for enterprise applications while they wind down Fibre Channel drive development.

The key reasons for this transition are that SAS drives have the same or better reliability and data availability characteristics of Fibre Channel drives but are less expensive. Indeed, not only the drives, but also the entire SAS infrastructure in a storage system is less expensive and, therefore, more powerful systems can be offered at a better price than storage systems with Fibre Channel drives. It's all about cost-effective performance.

When the Adaptable Modular Storage 2000 family was released, Hitachi Data Systems became the first in the industry to ship a storage system that scaled beyond 100 SAS drives. The newly released Hitachi Virtual Storage Platform is an enterprise class storage system that scales to more than 2,000 SAS drives. These storage solutions cover a wide range of requirements from medium sized businesses to the largest of enterprises and, with SAS technology, do so with the best price/performance ratios in their class.

Introduction: Market Projections for SAS Adoption

The hard disk drive (HDD) industry has been shipping enterprise disks with SAS interfaces since 2004. Initially, such disks were installed in servers and rapidly replaced the internal SCSI drives. Later, lower cost disk storage systems were introduced with SAS drives. Industry projections show that SAS disk shipments will grow faster than the market at the expense of Fibre Channel disks. Indeed, the disk vendors are planning to limit their development of new Fibre Channel disks in the next few years. SAS and SATA enterprise disk development is expected to continue into the foreseeable future with capacity increases and an interface roadmap that goes to 6Gb/sec and beyond. Fibre Channel disk development, however, will be capped at 600GB and a 4Gb/sec interface.

"Fibre Channel is dead. It's just a matter of time," said John Monroe, a research vice president in the storage group at Gartner, Inc., in reference to the native HDD interface. The key reasons for this transition are that SAS drives have reliability and data availability characteristics that are the same or better than Fibre Channel drives but are less expensive. Indeed, not only the drives, but also the entire SAS infrastructure in a storage system is less expensive. Therefore, more powerful systems can be offered at a better price than storage systems with Fibre Channel drives.

The forecasted shipments of enterprise disk drives by interface type are shown in Figure 1. Note the rapid increase in SAS drives over the next four years while Fibre Channel drive shipments are projected to cease in 2014.

Figure 1. Worldwide HDD Shipments for Enterprise Applications, 2009-2014 (All Form Factors by Interface)

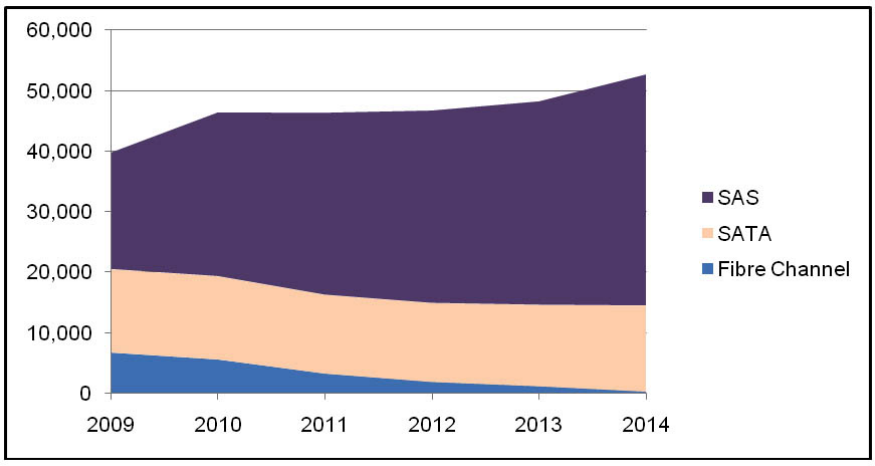

Source: IDC 2010

SAS Technology in the Hitachi Adaptable Modular Storage 2000 Family

The Hitachi Adaptable Modular Storage 2000 family uses a SAS interface to transport data from the controllers to the disk drive trays. The Adaptable Modular Storage 2000 family includes the AMS 2100, AMS 2300 and AMS 2500 systems, and it scales to 480 disk drives.

Each model in the family has a controller tray with dual controller boards. While the AMS 2500 controller tray does not have any drive slots, the AMS 2100 and AMS 2300 controller trays have slots for the installation of up to 15 internal disks in the controller tray. Additional storage can be added to the system by installing disk trays. Standard disk trays are 3U high and have 15 disk slots each while high density disk trays are 4U high and have 48 high capacity SATA disk slots.

Every controller board in a system controller tray has a DCTL RAID processor that connects to either one or two SAS I/O controller processors (IOC) depending on the model. The AMS 2100 and AMS 2300 have one IOC per controller board while the AMS 2500 has two. Each of the IOCs has two SAS multiplexor (SMX) ports that connect to a SAS expander (SXP) on one of the internal or external disk trays. The connection is made with a wide (x4) SAS cable that provides four 3Gb/sec SAS links. An illustration of the SAS IOC and its connections is shown in Figure 2.

Figure 2. Hitachi Adaptable Modular Storage 2100 and Adaptable Modular Storage 2300 SAS Design

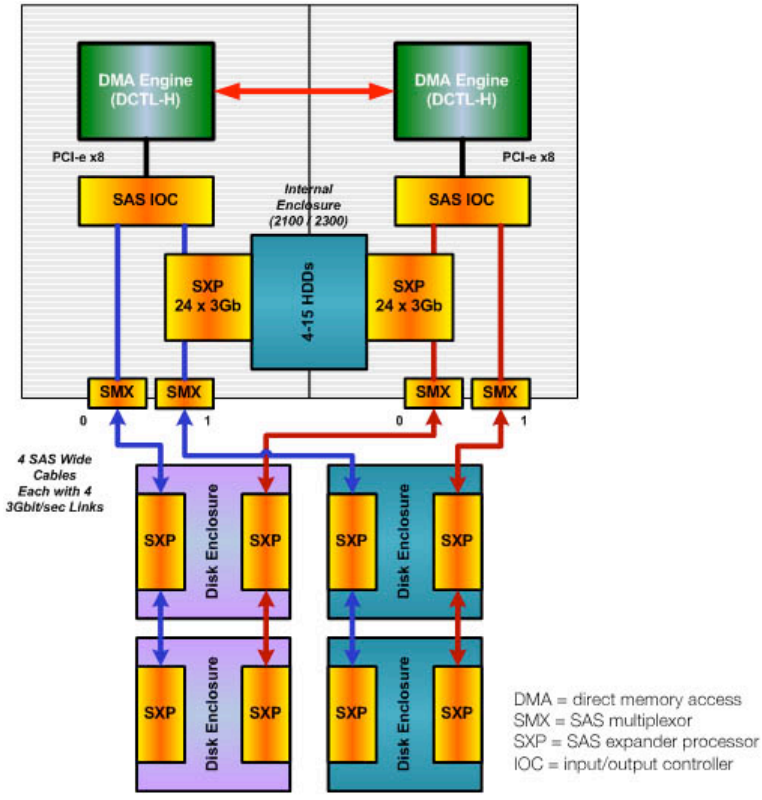

The SXP is a SAS expander processor and functions to establish the links between components. Two SXPs reside on each controller tray with internal disks (AMS 2100 and AMS 2300 systems) and each standard disk tray. There are four SXPs on each high density disk tray. On a standard disk tray, one SXP is connected to Controller-0 and the other is connected to Controller-1. As a result, each controller board has access to all of the disk trays. Each disk tray has eight SAS links for communications between the disks and the controllers.

Figure 3 illustrates how the SXPs access the disk drives in each standard tray. Since a SATA disk drive can be installed in a SAS slot, there is a common tray for SAS and SATA disks. Each SAS disk is dual ported so that either SXP can access the drive. The SATA disks are connected to an AA MUX (multiplexor) component, inside the disk canister, that allows a single port SATA disk to be connected to each SXP.

Figure 3. Standard Disk Enclosure

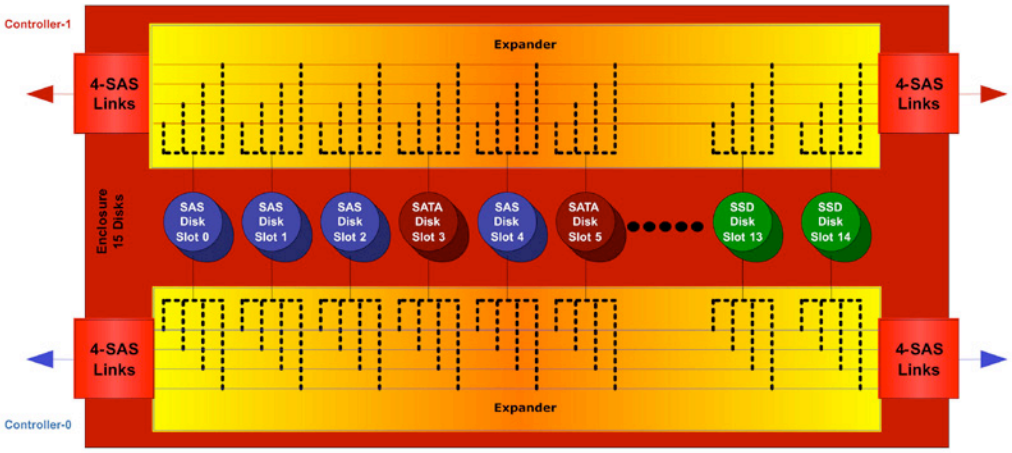

The high density disk tray is logically divided into two trays of 24 SATA drive slots each. Each half of 24 slots has two SXPs. Just like the standard tray, one of these SXPs is connected to Controller-0 and the other is connected to Controller-1. The same holds for the other half of the high density tray, as shown in Figure 4.

Figure 4. High Density SATA Disk Enclosure

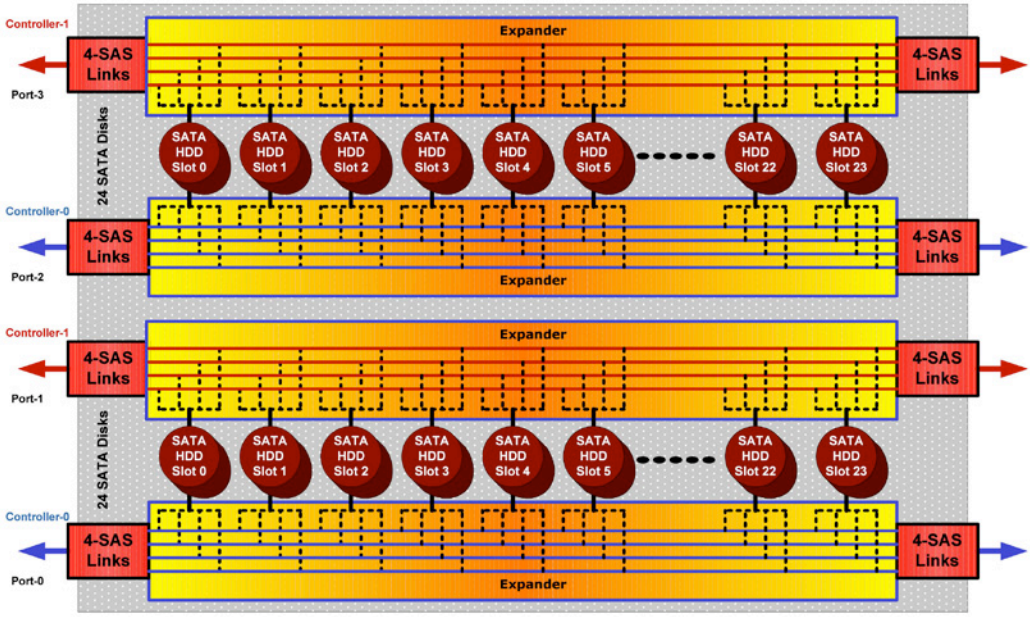

Whenever an I/O request is processed by the controller board, the SAS IOC will use one of its eight SAS links for communicating with a disk. Eight links are available on each of the controllers of the AMS 2100 and AMS 2300 systems and 16 links are available on each controller of the AMS 2500. Therefore, these systems have plenty of back-end bandwidth available. The SMX port on the controller will route the I/O to the first SXP expander. The SXP chip provides the actual connection between the IOC SAS links and the disks in an enclosure. Since the AMS 2100 and AMS 2300 controllers have internal disks, there are two SXP chips in their controller modules. There are no SXP chips in the AMS 2500 controller module since there are no internal disks.

If the target disk address resides on the same tray as the first SXP, then the expander will route the request from one of its four SAS links to the disk. Alternatively, if the target disk address is not one of the disks on the same tray as the SXP, then the I/O will be routed out the back of the expander to the next disk tray to which it is connected. This process continues until the SXP is reached that has a direct connection to the targeted disk.

SAS Technology in the Hitachi Virtual Storage Platform

The overall architecture of the Virtual Storage Platform (VSP) SAS Engine is similar to, though far more powerful than, the one used in the Adaptable Modular Storage 2000 family. It is the only enterprise-class storage platform with a SAS architecture.

Each VSP controller has a pair of back-end directors (BED) that has dual SAS Controllers (SPC). Each SPC has eight 6Gb/sec SAS links that are connected to the hard disk units (HDU) that are

installed in a disk cabinet (DKU). The VSP has eight 6Gb/sec SAS links per HDU. The switches are able to connect eight drives within that HDU at the same time, or to bypass those drives on some links and go to the next HDU.

This provides 32 6Gb/sec SAS links (16 2W cables) to up to 640 LFF (large form factor, 3.5in.) disks or 1280 SFF (small form factor, 2.5in.) disks. An interconnected, dual chassis system would have two of these structures with up to 1280 LFF drives or 2048 SFF drives. The design of a single chassis VSP and the SAS links are illustrated in Figure 5.

The AMS 2000 family systems and the VSP are both similar in that I/Os can go through any host port on the controller and still reach their SAS controller without performance degradation. This symmetrical aspect mitigates path management complexities.

Figure 5. Back-end SAS Link Layout (High Performance Configuration)

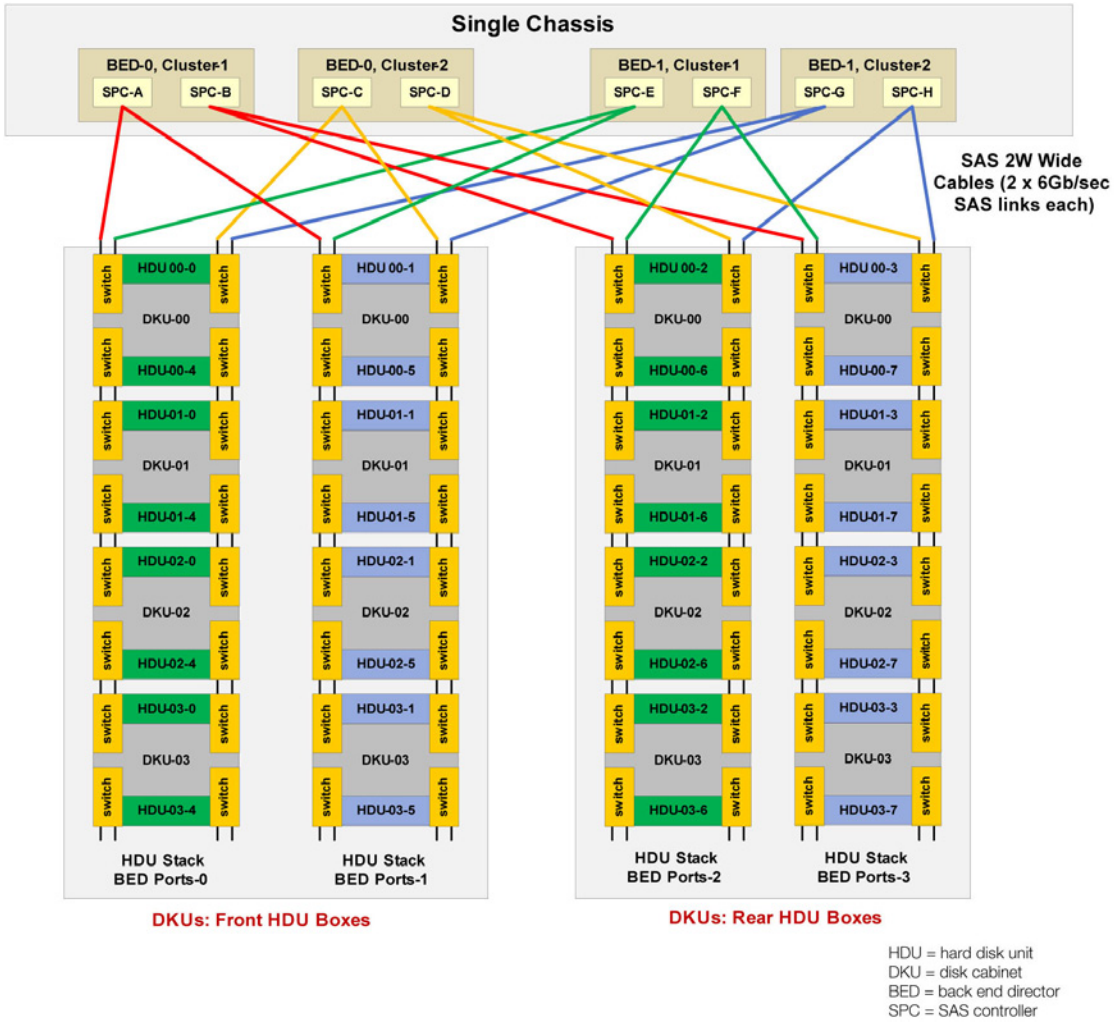

Each HDU has two drive expander switches. On the LFF HDU, there are slots for 10 3.5in. drives (2TB SATA and 400GB solid state disk or SSD). On the SFF type of HDU there are 16 drive slots (all SAS drives and the 200GB SSD). All eight of the HDUs per DKU are of the same form factor. Figure 6 provides an illustration of the SFF HDU while Figure 7 shows the LFF HDU. Each SPC processor determines which of its two links will be used to connect to each drive.

Figure 6. View of the Two SAS Switches in a SFF HDU

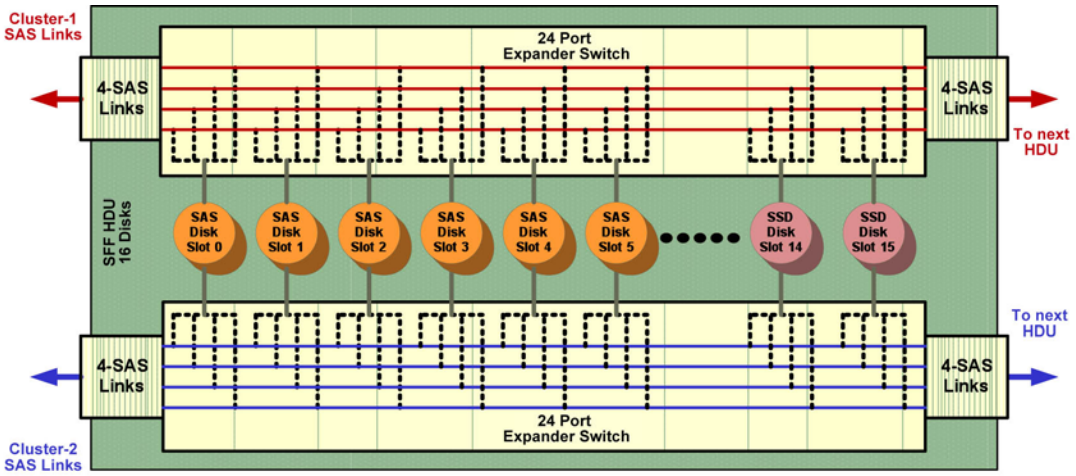

SFF HDU = small form factor hard disk unit

Figure 7. View of the Two SAS Switches in an LFF HDU

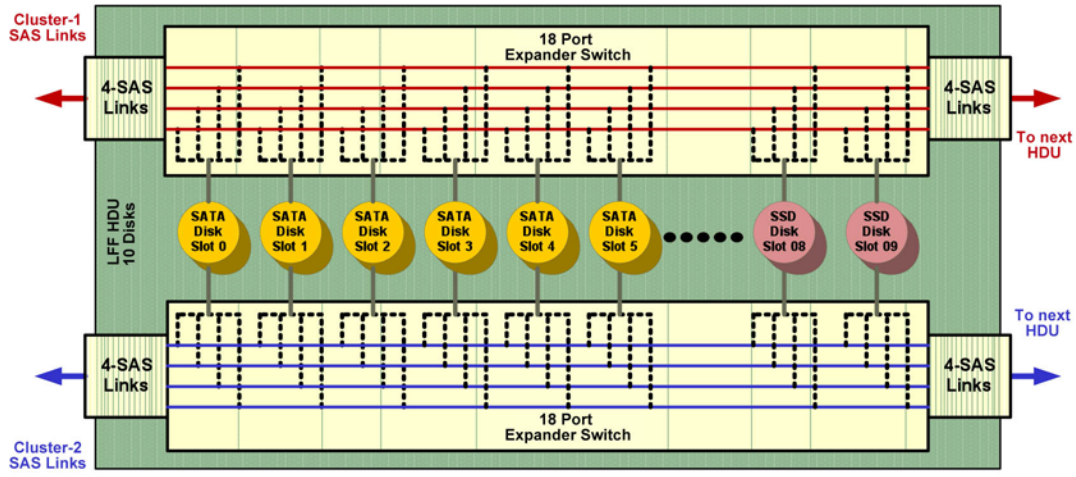

SFF HDU = small form factor hard disk unit

Benefits of SAS

Three main business benefits are driving the transition from Fibre Channel to SAS drives.

Cost-effective Performance

The recently released Hitachi Virtual Storage Platform is an enterprise-class storage system with twice the number of back-end data links as the prior generation Hitachi Universal Storage Platform® V (USP V). The speed of each link has been increased from 4Gb/sec Fibre Channel to 6Gb/sec SAS. With more and faster links available to the drives, the VSP has outstanding performance characteristics.

Storage system performance is especially important for customers who install flash drives or solid state disks in their storage systems. These customers want the highest I/O possible for selected applications. Storage systems that have a small number of Fibre Channel data links to the flash drives will quickly saturate their back end. Therefore, the high number of I/Os that flash drives deliver cannot be fully realized by the systems. With 4 times the number of data links, in the case of the AMS 2000 family or 1.5 times the link speed, in the case of the VSP, these storage systems are designed to better deliver the blazing performance that flash drives provide.

The older Adaptable Modular Storage 2000 family is a midrange storage design that has a 3Gb/sec SAS back end. While a single link on an AMS 2000 family system does not have the same throughput as a system with a 4Gb/sec Fibre Channel back end, one must evaluate the performance capabilities of the entire system. In that regard, the AMS 2000 family systems are superior for the following reasons:

a. Systems in the Adaptable Modular Storage 2000 family have more back-end links than any competitive modular system, allowing for more concurrent I/Os. AMS 2100 and AMS 2300 systems have 16 SAS links and the AMS 2500 system has 32 SAS links. Comparable Fibre Channel-based modular systems typically have four or eight loops.

b. SAS-based switch architectures are not impacted by loop arbitration wait times that slow down FC-AL (Fibre Channel-Arbitrated Loop) based systems that share a common loop.

c. SAS is always full duplex, which means that any device can transmit and receive I/Os simultaneously.

The high level of performance that the Adaptable Modular Storage 2000 family achieves has been validated by the Storage Performance Council (SPC) benchmark results, specifically SPC Benchmark-1™ (SPC-1). In SPC-1 benchmark testing, the Adaptable Modular Storage 2500 achieved the fastest throughput results among all midrange storage competitors with dual controllers. It delivered an impressive throughput result of 89,491.81 SPC-1 IOPS and an 8.98 millisecond average response time. The price/performance ratio of these results was a very low US$6.71/SPC-1 IOPS. Only the Adaptable Modular Storage 2100, at $5.95/SPC-1 IOPS, has a better price/performance ratio amongst modular systems. In fact, the three models have the best price/performance ratios for modular systems as measured by SPC-1.

Since the traditional methods of achieving required levels of performance are either to spend more money for more powerful systems, or to spend more time managing the systems for performance optimization, the VSP and the AMS 2000 family systems can save both capital and operational expenses.

Investment Protection

Most organizations that are buying new midrange or enterprise storage systems today plan to operate their units up to seven years. Over the projected lifespan, capacity upgrades will be needed to keep up with growth requirements. Disk drives must be available for purchase to manage data capacity growth. As we see from the IDC projections, the availability of Fibre Channel disk drives after 2014 is in question which puts the long-term lifespan of Fibre Channel-based storage systems at risk. Without the ability to add disks for expansion, these systems may be prematurely downgraded or even de-commissioned. Storage systems based on SAS technology offer better investment protection since the long-term availability of drives is not in question.

Flexibility

The SAS and SATA drive connectors allow both types to be intermixed in the same drive enclosure, eliminating the need to purchase unique disk trays and simplifying cabling. SATA commands are encapsulated inside SAS packets using the SATA Tunneling Protocol of the SAS specification. Technologies such as a connector dongle or FATA allow Fibre Channel and SATA drives to be installed in the same tray. However, this requires that the I/O to the SATA drives be converted from Fibre Channel to SATA. This translation involves additional expense and performance robbing overhead.

SAS Reliability

SAS drives have been used for many years in data center environments as the internal drives that store the operating system in servers. These environments require levels of reliability that can be business critical. SAS and Fibre Channel disks are nearly identical mechanically. They have the same heads, disks and motors, and only differ in their interface electronics and connectors. As a result, they have the same reliability ratings. Their meantime between failure (MTBF) and duty cycle ratings are identical.

The SAS interface has an advantage in how it handles error detection. If a component on a SAS link fails, not only will that component be mapped out, but the failure also will be broadcasted according to the SAS specification and reported by the system monitoring interface. Compare this to how failures are handled on a Fibre Channel loop. With Fibre Channel, the failure is reported at the loop level and not at the component level. Component-level reporting allows failures to be corrected in less time and with improved accuracy.

Conclusion

With the introduction of the Adaptable Modular Storage 2000 family and the Virtual Storage Platform systems, Hitachi has been a leader in the inevitable transition of Fibre Channel to SAS disk interface for high performance storage. SAS offers the same or better reliability and data availability characteristics. With a SAS-based architecture, the Adaptable Modular Storage 2500 was able to achieve the best-in-class performance results as measured by the SPC-1 benchmark. Hitachi continues to lead the way in storage with high performing systems that are priced for any organization's budget.

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com