IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Addressing Business Challenges with A Scalable Storage Infrastructure

Sponsored by: HDS

Richard L. Villars
October 2008

EXECUTIVE SUMMARY

In today's rapidly changing business environment, companies rely on an expanding set of applications to compete, and these new applications have significant effect on how they organize their datacenters and store structured and unstructured information. One major consequence of these changes is the emergence of modular storage systems as the "primary" storage platform, in terms of capacity deployed and IT staff resources required.

A number of IT technologies and data center practices are accelerating this development, changing both the scope and type of modular storage systems that organizations need to deploy. The most influential new developments are:

- Use of virtualization to improve asset use and application availability
- Replication of information to boost recoverability and meet retention mandates
- Accelerated creation and long term archiving of rich digital content

New modular storage systems like those in Hitachi's AMS 2000 family are addressing these new requirements by delivering:

- Superior throughput, resiliency, and processing capabilities through use of flexible and dynamic active/active controller architectures
- Improved data access rates, reliability, and energy consumption through use of new disk storage technologies such as Serial Attached SCSI (SAS)
- Optimal in storage capacity and energy efficiency through use of disk spin down technologies to better match the needs of large backup and file repositories
- Logical abstraction between the physical storage elements and the data volumes that virtualized servers and applications need to access

Hitachi new systems and supporting software solutions address many IT organizations most critical information management needs. Companies must, however, also look for a business partner who can help them select, integrate, and deploy these new storage systems as part of a complete solution that meets their unique requirements.

THE CHALLENGE OF INFORMATION MANAGEMENT IN TODAY'S ENTERPRISE

Today, your company relies on an expanding set of applications to compete in a rapidly changing business environment. You rely on email, ecommerce systems, and Web sites to conduct business with customers and business partners. You are collecting, storing, and analyzing more information about products, customers, and transactions. You are digitizing records, images, and other unstructured data to offer new services and comply with evolving government regulations.

The expansion in the range of applications that companies are developing and deploying has a significant effect on how they organize their datacenters and store structured and unstructured information. IT executives must judge all storage offerings based on three business requirements. Does the system

- Enable more effective use of all types of information by a wider range of users
- Reduce and control the cost of doing business by boosting asset utilization
- Ensure the integrity of the business and its information assets in the face of natural disaster, systems failures, or outside regulatory oversight

Key Challenges Reshaping Storage Requirements

Prior to the mid-1990s, storage was a primarily a subsystem (e.g., a hard disk drive or tape drive) within a system (server or PC). With data intense and mission critical applications running on large servers. As a result, expanding storage capacities, boosting IOP performance, and improving data integrity became critical new storage requirements. Developers of the first generation of high-end storage systems sought to provide more capacity/higher reliability (e.g., RAID) and broader accessibility (e.g., storage area networks) for structured data in transaction-intensive applications.

In the late 1990s, storage systems suppliers began to offer modular storage systems, based on dual controller architectures, to better address the cost and scale needs of departments and medium-sized businesses. Very quickly, however, companies of all sizes recognized that modular storage systems were becoming their "primary" storage platforms in terms of capacity deployed and IT staff resources required to manage storage systems.

Today, a number of IT technologies and data center practices are accelerating this evolution and changing both the scope and type of modular storage systems that organizations need to deploy. The most influential new developments are:

- The adoption of virtualized IT technologies within the datacenter to improve asset use and improve application availability
- The replication of information (both local and remote) to improve recoverability, enable business analytics, and meet data retention/discovery mandates
- The explosion in the creation and archiving of unstructured data in businesses

Virtualization's Unintended Consequences for Modular Storage

Server virtualization is one of the most important datacenter trends of the past several years. The ability to virtualize servers and reclaim excess capacity is of interest to datacenter managers in all companies. These virtualized servers are much more likely to access networked modular storage, a prerequisite for companies that want to use advanced features such as workload migration and disaster recovery.

Storage assets allocated to virtualized servers can deliver significant value, but limitations in the active/passive architecture of many traditional modular storage systems often leads to "unintended consequences" that quickly minimize or overwhelm the original benefits associated with server virtualization. These unintended consequences include:

- Performance imbalances between virtualized servers, storage networks, and back-end modular storage systems that lead to long and expensive network re-architecting and storage migration projects
- Struggles by storage administrators to match the rapid pace of provisioning associated with virtualized servers in their own provisioning of storage resources
- Significant storage "over-provisioning" as virtual machines (VMs) proliferate and force change in system backup and recovery practices.

Beyond, these technical consequences, adoption of virtualized servers often changes an organization's approach to all IT investments. Does the investment improve asset utilization efficiency (in terms of capital investment and ongoing operations), enhance IT responsiveness to business requirements, or boost service reliability?.

Boosting Data Availability and Value with Storage Replication

Customers and governmental organizations expect companies to quickly deliver the right information to the right person in a timely fashion, so control of information assets is critical. Information must be protected from permanent loss or extended unavailability and that information is easy to rediscover and reanalyze. The capabilities required in storage systems to meet these goals go well beyond those found in traditional modular storage systems. In the past, they required investments in expensive, higher-end storage and/or expensive replication software

IT executives want modular solutions that provide high levels of consolidated data protection while balancing security, cost, and business requirements. These availability, accessibility, and analytic requirements are spurring an explosion of options in the replication of data. Vendors are deploying an increasingly diverse array of storage systems designed to meet different roles (e.g., disk-based backup/recovery, email/records retention/archive, Business analytics/e-discovery, and virtual image management for virtual servers and desktops). This proliferation of systems storing replicated data has a number of unintended consequences:

- An over-duplication of data leading to reduced asset utilization and rising costs
- A potential loss of centralized control over the distribution and eventual destruction of data

Organizations' Evolving Content Requirements

Transactional applications such as ERP, CRM, and OLTP that typically create and access structured data remain major consumers of new storage capacity; however, new applications that generate and access unstructured data are now more voracious consumers of storage capacity in many enterprises (See Figure 1). These include:

- Digital image and video archiving systems in healthcare, entertainment, security , and product design/manufacturing
- Collaborative applications such as email and SharePoint
- Virtual machine images (for servers and desktops) that enable improve application availability, system recovery, and rapid provisioning

The continued expansion in unstructured information is already having a dramatic impact on organizations' storage environments, with IT executives coming to the realization that managing such file-base data will become the primary task for many storage administrators. These administrators need storage solutions that can support reliable and faster access to significantly more and larger files; ensure that the sharing of business content is more cost-efficient and secure; and enhance (aka monetize) the value of unstructured data.

FIGURE 1

Worldwide Enterprise Disk Storage Consumption Model, Capacity Shipment Share by Segment, 2007-2012

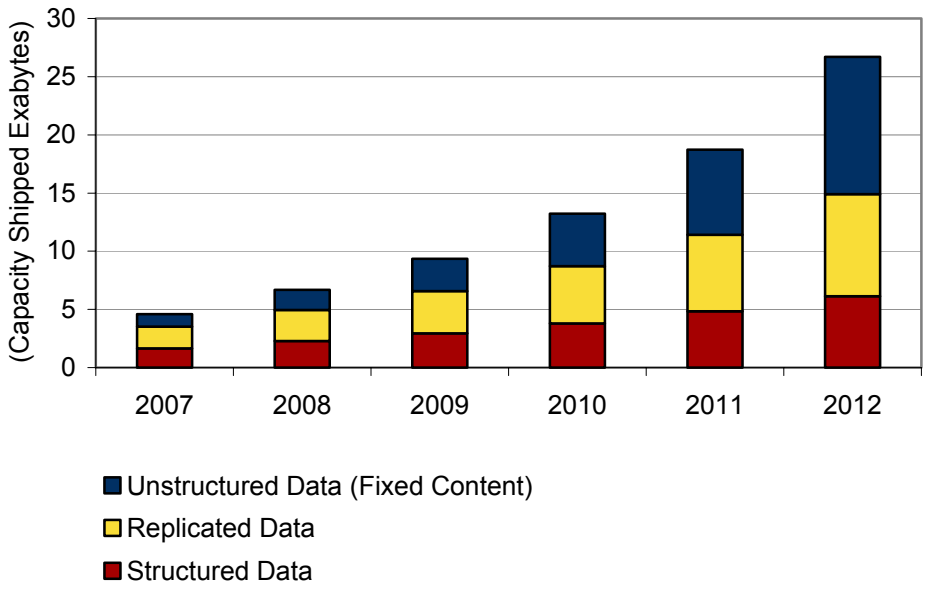

Source: IDC, 2008

CHANGING ROLES AND REQUIREMENTS FOR MODULAR STORAGE SOLUTIONS

The shift to a virtualized IT environment, expanded data replication, and richer digital content, will drive future developments in modular storage systems and supporting storage software. New systems must directly address the fast-expanding and increasingly diverse needs of organizations through a number of technology initiatives. The most significant in the coming year are:

- Improvements in modular storage system throughput, resiliency, and processing capabilities through use of flexible and dynamic active/active controller architectures and advanced server processors
- Improvements in data access rates, reliability, and energy consumption through the adoption of new disk technologies such as Serial Attached SCSI (SAS)
- Improvements in storage capacity and energy efficiency through deployment of high capacity disk storage technologies (e.g., SATA) and techniques such as disk spin down to better match the needs of large backup and file repositories
- Introduction of logical abstraction (e.g., block level virtualization, thin provisioning, and data-deduplication) between the physical storage elements and the data volumes that virtualized servers and applications need to access

These technology changes will have positive effects on operational efficiency, scalability, and responsiveness to business requirements; however, they are also likely to lead to significant disruptions in storage hardware design, storage management processes, and storage administrator responsibilities. IT executives need a storage systems supplier that can deliver the right combination of scalable modular systems and provide the right supporting software and services to make modular storage and integrated part of the data center, not an impediment to growth.

HITACHI'S MIDRANGE OFFERINGS: SCALEABLE STORAGE FOR BUSINESS

Hitachi Data Systems (HDS) is a wholly owned subsidiary of Hitachi, Ltd. (NYSE: HIT) and is a leading supplier of storage systems, software and services to enterprises around the globe. HDS develops a variety of storage hardware and software products designed to address large and mid-sized companies' growing and evolving information requirements.

The remainder of this white paper examine the latest modular storage offerings from Hitachi and assesses their ability to help organizations control costs, sustain innovation, and ensure ongoing business operations. Such solutions are designed to:

- Deliver more scalable, flexible, and reliable IOP performance
- Provide simple yet scalable options that support ongoing capacity growth while boosting operational and energy efficiency

- Centralize the process of organizing and managing all storage assets
- Address the changing needs of midrange customers with flexibility in platform choice and unified management of storage systems

Hitachi's AMS 2000 Modular Storage Solution

In 2008, HDS is delivering a new generation of modular storage solutions designed to address the specific needs of mid-sized organizations. These include a new generation of modular storage systems, the 2100, 2300, and 2500 which are designed to be more scaleable and flexible while also being simpler to configure and deploy than traditional modular storage systems.

These systems allow companies to deploy multiple tiers of disks based on size, performance and cost within a single, scalable system with maximum system sizes ranging from 47 TB (All SAS on a 2100) to 472 TB (SATA on a 2500). The systems also support anywhere from 4 to 16 Fibre Channel (FC) or 1 Gbps iSCSI SAN connections and from 4 GB to 32 GB of system cache. Key elements in the AMS 2000 products include:

- Implementation of a Dynamic Load Balancing Active/Active controller architecture for optimal reliability and IOP performance in support of virtualized IT
- Use of a flexible disk architecture (SAS/SATA intermix) across the product line to provide improve system availability, backend performance, and energy efficiency
- Optimized modern architecture for business-critical applications that require performance and availability, and a superb fit in virtual server environments.

Hitachi specifically developed these modular storage systems, as well as complimentary products in its midrange portfolio for storage virtualization, advanced replication, data migration and consolidation (Hitachi USP VM) and file-based storage (Hitachi High-performance NAS Platform), to address the performance, data and management intense requirements of virtualized servers, collaborative applications (email), and large content repositories.

Delivering Availability and Longevity With SAS

A key component in Hitachi's new AMS 2000 product line is the adoption of backend SAS architecture for all systems. Historically, modular storage systems relied on a Fibre Channel backend which called for use of shared network resources, FC disk drives and a SATA to FC bridge for SATA support.

Hitachi is a leader in leveraging SAS disk and connectivity technologies (which are already widely deployed in internal servers), in modular external storage systems. This new architecture allows Hitachi to deliver modular storage solution that:

- Supports significantly faster data transfer rates (up to 9600 MB/sec of total backend bandwidth

- Allows for the intermix of SAS and SATA drives on the same tray (an option that is not practical in many FC/SATA environments)

- Enable much faster disk failure detection and isolation (significantly improving system and data availability)

Boosting Performance/Availability With Dynamic Load Balancing

Traditional modular storage systems support dual controller configurations but require administrators to assign individual LUNs to a specific controller. In cases where a controller fails, organizations had the choice either losing access to that controller's assigned LUNs until repair, or pre-configuring an alternate path which required additional server software. This forced administrators to perform time-consuming configuration settings for each LUN, and could exact a significant performance penalty on the entire system. This limitation also makes it necessary to shut down access to part of the system whenever microcode is updated (See Figure 2).

FIGURE 2

Updating Microcode in Tradition Modular Storage

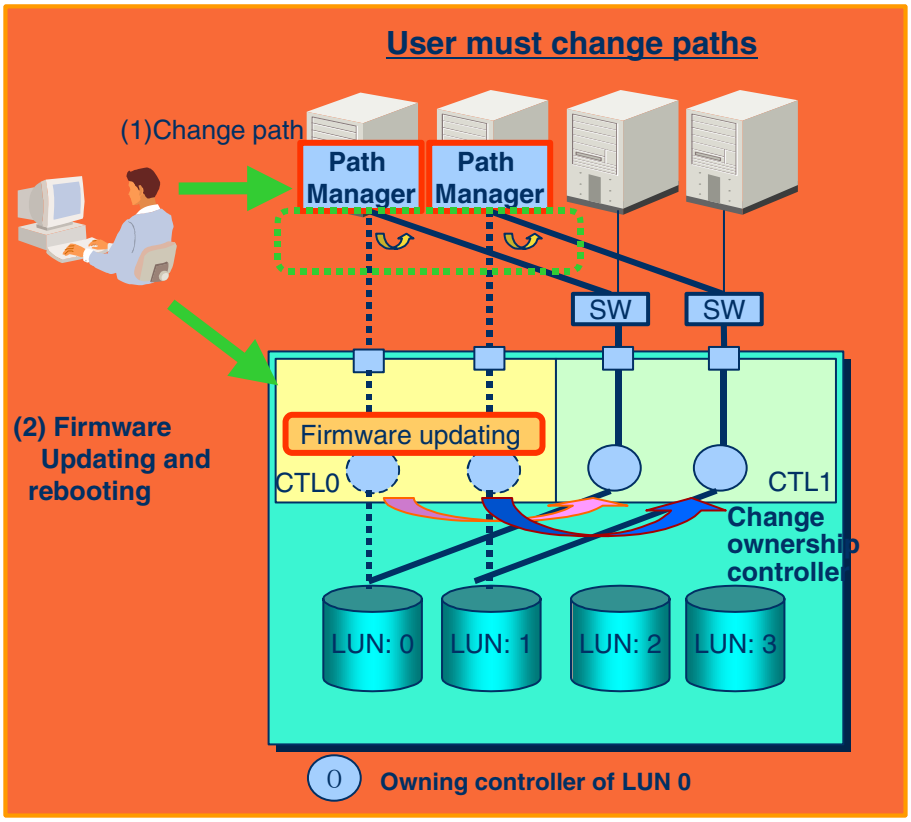

Source: HDS, 2008

The growing connectivity of virtualized servers running many different applications is now exposing another shortcoming of this approach. Virtualized environments make it much harder to predict and balance IOP loads across LUNs, so administrators are unable to easily reduce "hot spots" in data demand. This shortcoming makes balancing systems performance difficult and reduces effective utilization rates.

The AMS 2000 product line includes and active/active controller architecture to address many of these concerns. It then went a step further and enables a feature which Hitachi calls the Hitachi Dynamic Load Balancing Controller architecture (See Figure 3). This architecture:

- Eliminates hot spot and/or utilization imbalance issues
- Provides resiliency for planned and unplanned outages as data traffic is non-disruptively moved from a non-operational to an operational controller
- Eliminates mandatory deployment of proprietary path management software and preferred paths so any host can access any LUN without a performance penalty
- Enables non-disruptive microcode updates with little impact on operations

FIGURE 3

Benefits of Hitachi's Dynamic Load Balancing Controller Architecture

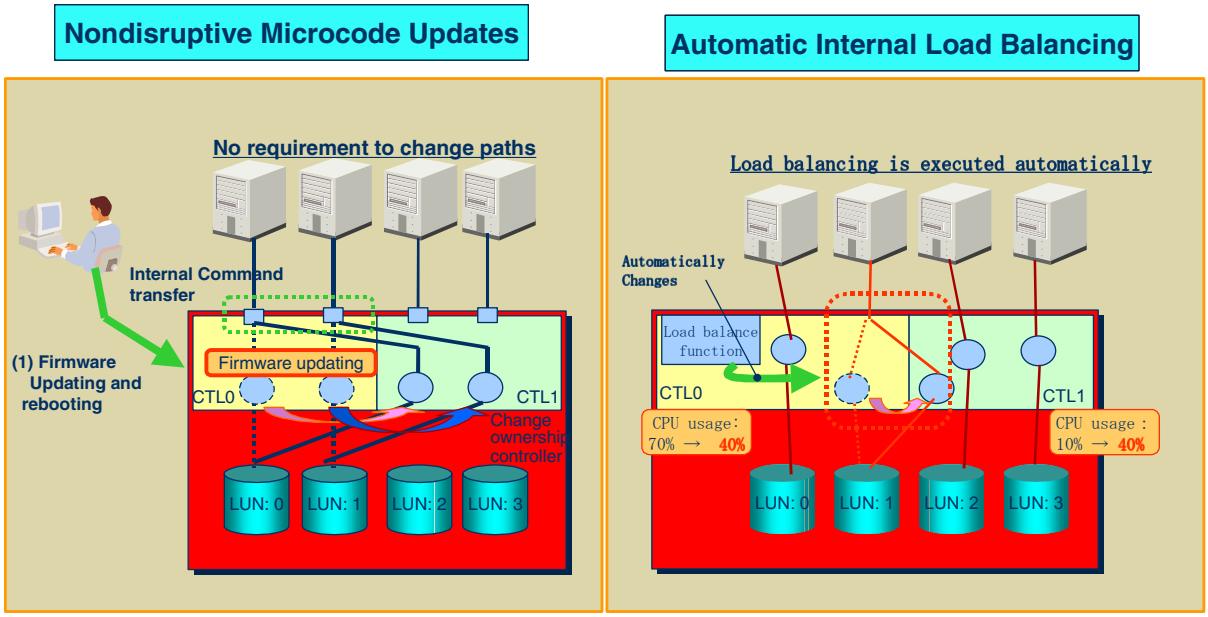

Source: HDS, 2008

Delivering More Efficient Modular Storage

As noted above, the benefits that organizations are obtaining from server virtualization are influencing their evaluation of other IT environments. In storage, this phenomena is most apparent in IT executives demands for more efficient storage solutions. This efficiency demand has two major components, boosting the actual capacity utilization rate of purchased storage systems and reducing the energy consumed by those systems.

This focus goes beyond just substituting lower cost secondary storage (SATA) for higher cost primary storage in a tiered storage environment . It also means that IT executives want to reduce "over provisioning" and reduce power consumption for disks that contain infrequently accessed data. Hitachi's goal is to address both of these requirements in its modular storage environments through the introduction of dynamic provisioning and support for intelligent disk management

Hitachi Dynamic Provisioning

In 2007, Hitachi added thin provisioning capabilities on its high end USP storage product line with the introduction of Hitachi Dynamic Provisioning. This allowed large enterprises to boost the utilization levels for both primary and secondary storage. Later that same year, Hitachi added support for this capability to its midrange portfolio with the USP VM networked controller which delivers similar capabilities for pools of modular storage systems such as the AMS 2000.

With the introduction of the AMS 2000 product line, Hitachi is committed to delivering dynamic provisioning on individual modular systems. IT staff can:

- Improve storage capacity utilization, minimize upfront costs, and reduce power consumption by allowing virtual disk storage to be allocated based on future application requirements without deploying unneeded physical capacity
- Purchased and deploy physical capacity at a later time (and at a lower cost) when actually needed, without any disruption to mission-critical applications
- Leverage the active/active controller and the SAS backplane-enabled performance of the AMS 2000 in a tiered storage/virtualized environment with USP-VM for example

Intelligent Disk and System Operations

Beyond the indirect power savings offered through thin provisioning (one doesn't need power for disks that aren't actually installed), Hitachi also added a number of direct energy management capabilities to the AMS 2000. These include:

- The ability to spin down drives in RAID groups that are infrequently accessed
- Support for SATA drive features that enable multiple power modes and the ability to park heads when idle for more than 2 hours (e.g., when supporting VTL)
- Array-level monitoring of system cooling requirements and adjustable fan speeds

These solutions allow the AMS 2000 products to respond to increasingly variable operating conditions without sacrificing energy efficiency.

Hitachi Storage Navigator Modular 2

The final component in Hitachi's modular storage portfolio is a set of software services that provide advanced management services for the all of Hitachi's AMS platforms.

Hitachi's mature portfolio of storage and data management software solutions, the **Storage Command 6.0** suite, includes storage area management modules that provide common management capabilities and an intuitive user interface for a heterogeneous storage infrastructure. The HiCommand Storage Services Manager acts as the main console for Hitachi's heterogeneous storage infrastructure management software, providing SAN visualization and reporting, asset management, performance and capacity monitoring and planning, and policy-driven event management.

Recognizing that many modular storage customers don't want to deploy such management platforms, but still need management on their installed systems, Hitachi Storage Navigator Modular 2 can also provide simple, wizard-based system configuration facilities plus a number of advanced capabilities:

- Host and LUN Management (including LUN grow/shrink and RAID expansion)
- Cache Partitioning and Cache Residency (tuning cache to meet different workloads)
- Replication setup and management (enabling improved data protection/recovery)

Challenges for Hitachi

Demands or modular storage systems will continue to evolve as new environments (e.g,, virtual desktops) and new use cases (e.g., business analytic warehouses) emerge in the coming years. The development of new technologies and product lines such as the AMS 2000 and the USP VM will make it easier for IT managers to adapt storage systems for consolidation, archiving, compliance, and business continuity.

Having a broad catalog of products that meet demands for better cost performance, availability, and reliability isn't enough, however. IT managers need their storage suppliers and value-added resellers to provide simple solutions that quickly address specific business requirements and work seamlessly with other technologies such as server virtualization.

Hitachi is broadening its mid-range options but addressing the shifting needs of IT organizations ultimately takes more than new hardware and software. IT managers need a business partner who can help them select, integrate, and deploy the right combination of hardware and software products to deliver complete solutions that meet their unique requirements.

Over the past three years, Hitachi has made significant investments in developing strong business partners around the globe. These partners focus on large enterprises with advanced storage requirements and mid-sized enterprises that have different needs and expectations. With this next generation of modular storage solutions that bring advanced capabilities to new environments, Hitachi must continue to extend the capabilities of its business partners to target new applications and deliver more effective implementation and support services.

FINAL THOUGHTS AND ESSENTIAL GUIDANCE

HDS, with its midrange, modular storage solutions, is developing specific solutions that target the unique challenges of virtualized IT, disk-based data protection, and large content repositories. It is delivering a set of targeted solutions that support IT consolidation through the use of virtualization, enables cost-effective long-term data archiving, and addresses the specific storage management requirements of IT managers that recognize modular storage is their primary storage environment.

You as an IT manager must look beyond products. You need partners that have the forethought to incorporate management into both the fundamental product design and into the deployment process. Ongoing administration/management is the most overlooked and misunderstood issue in most IT deployment plans. Your storage solutions partner and your implementation partner must provide the tools and service to help you develop a sound management and provisioning process that meets both today's and tomorrow's needs.