

Services Oriented Storage Solutions

White Paper

By Hu Yoshida and Peter Smails

May 2007



Contents

The Promise of Storage Area Networks	2
The Roots of Services Oriented Storage Solutions	2
Virtualization, the Key to Services Oriented Storage Solutions	3
Extending Services Oriented Storage Solutions to Storage Application Servers	4
Proof Point for Services Oriented Storage Solutions.....	4
Platform Summary	5
Services Oriented Storage Solutions in Action	6
Storage Management	6
Tiered Storage	6
Network Attached Storage (NAS).....	6
Business Continuity	7
Active Archiving.....	7
Solutions for Smaller Organizations	8
Services Oriented Storage Solutions Summary	8
Appendix A: More Information	9

Services Oriented Storage Solutions

White Paper

Connecting IT with business has been the mantra of IT organizations for many years. However, the reality often finds the data center mired in redundancies and the proliferation of monolithic architectures and infrastructure. Progressive IT organizations have adopted a services-oriented approach to managing core IT functions. Services are increasingly defined in user's terminology and the IT infrastructure needed to support those services is mapped and managed to service level agreements (SLAs). Leading IT organizations are focused on dynamically provisioning and re-provisioning the entirety of the infrastructure (networks, servers, and storage) using policy-based engines to ensure that "Gold" level services always meet their SLAs. The ability to do this in a cost-effective manner is the trick. In the past, storage systems lagged behind servers and networks whose management tools have adapted to these needs. Hitachi Data Systems has changed all of that.

Services Oriented Storage Solutions apply service-oriented architecture (SOA) concepts to storage to deliver a platform that can be readily reconfigured and optimized to changing business requirements; these solutions deliver a process-oriented service approach to storage rather than the piecemeal, task-oriented approach, which leads to needless redundancies, over-subscription of storage, management complexity, and compliance exposure.

Hitachi Data Systems has been developing its storage platform with a services-oriented approach for many years. Some of the hallmarks of the Hitachi Data Systems strategy have been:

- Control unit virtualization with enhanced storage services that enable heterogeneous storage systems to interact and work in concert to optimize storage performance, data protection, and system availability
- An integrated portfolio of storage management, business continuity, and data mobility services that enable organizations to leverage a single set of tools for all of their storage and data management challenges

Building on these proven, enterprise-class capabilities, Hitachi Data Systems has expanded their strategy to include file and object services, enabling organizations to leverage a single platform for all their data storage requirements. Most importantly, these organizations can now respond more quickly to business and technology change and:

- Reduce cost and increase efficiency by reducing the complexity of their infrastructure and automating the process of storage management
- Boost utilization and reduce over-subscription of storage resources
- Cost-effectively address a growing array of structured and unstructured data types and applications
- Improve availability, reliability, and SLA consistency for midrange and small enterprise data applications
- Provide metrics and enable policies to measure and automate the use of storage services

The Promise of Storage Area Networks

Fibre Channel storage area networks (SANs) were introduced in the late 1990s to eliminate the islands of direct attached storage. SANs promised to consolidate storage and centralize management of storage and data resources. The assumption was that this consolidation and centralization would lead to better utilization of storage capacity, which was then about 20 percent to 30 percent. Utilization at that time really referred to allocation of storage across servers and applications. There were tremendous improvements in efficiencies around allocation of storage across direct attached servers, and SANs were successful in addressing them. SANs enabled more servers to attach to larger and larger capacity storage systems, and Storage Area Management tools were developed to map the resources in a large heterogeneous SAN. They did their job in networking servers to storage.


However, when we look at the utilization of storage today, more than 10 years after the introduction of SANs, utilization in many cases is still around 20 percent to 30 percent. What happened to the promise of SAN to increase utilization? There are many reasons for this low utilization. The primary reason is an over-subscription or over-buying of storage. As disk capacity prices declined, many storage vendors began to build larger and larger capacity storage systems to recover their costs, requiring customers to buy more than they need to get the latest functionality. Business units who do not want to share storage systems with other users cause excess storage capacity to be stranded. Niche storage solutions like archive or virtual tape libraries (VTLs) create new islands of storage that cannot be shared with other applications. Users who want to avoid the disruption of running out of space request more allocation than they will ever use, and then make copies of that allocation for backup, business continuance, development test, data mining, data distribution, etc. As the capacity of storage systems escalate, the effort to consolidate storage or migrate to more efficient storage configurations becomes too disruptive and more of the IT budget is spent on maintaining old technology than investing in new. When budgets are flat and data requirements explode it appears cheaper to throw storage capacity at the problem than hire the people or acquire the tools to manage it. Tony Asaro of Enterprise Strategy Group identifies this as a failure of SANs to network storage. He points out that SANs provide a network of host servers to storage but do not network storage systems to other storage systems...“aside from being on the same network, individual storage systems do not, in any way, interact with each other. They do not work in concert, but rather separately: individually; discretely; and therefore inefficiently.”

If storage systems could work together to share their capacity between applications while providing safe multitenancy and quality of service, if the same storage platform could support niche applications like archive or virtual tape, if the capacity could be allocated as it was used, if consolidation and migration could be done without disruption to the applications, and if we could use a common set of tools to manage growth, then we could address the problem of low utilization.

The Roots of Services Oriented Storage Solutions

Recognizing this problem, Hitachi introduced a storage controller with the ability to network storage resources within a storage system. In 2000, Hitachi delivered the Lightning 9900™ Series enterprise storage systems, which introduced the ability to dynamically change storage configurations within global cache through the introduction of a cross bar switch architecture and separation of control data from user data. Logical Disk Volumes within the storage array could work in concert, to deliver optimal performance and availability to host applications.

In 2002, the Hitachi Lightning 9900 V Series introduced the addition of virtual storage ports, which increased the connectivity of each Fibre Channel storage port to 128 virtual ports. Each had its own dedicated address space, so that applications that shared the same storage port could be secure from data leakage or data corruption from each other. Virtual porting helped with allocation rates across server types and applications. It complemented SANs by allowing more of the storage system to be shared unlike other storage systems with a limited number of



statically configured storage ports. Multiple applications, residing on different platforms like mainframes, UNIX, Linux, Microsoft® Windows, and Novell, could share the same storage system, and use different RAID protection, different performance disks, and set different priorities on port access. Not only were the Lightning 9900 V Series systems able to optimize storage for multiple applications simultaneously, they could also dynamically change the storage configuration as the requirements of the application changed. Hitachi also introduced a suite of management tools that could look into a heterogeneous SAN and discover, visualize, monitor utilization, and charge back usage of resources from an Oracle table space through host bus adapters (HBAs), switch ports, and storage ports to the Logical Unit. This gave customers visibility into their storage environment and enabled them to massively consolidate and optimize their infrastructure while greatly reducing cost, complexity, and management.

Virtualization, the Key to Services Oriented Storage Solutions

While the Lightning 9900 V Series systems were able to network storage in Hitachi storage systems, it was apparent that networking of storage resources had to extend beyond individual storage systems. In 2004, Hitachi introduced the Universal Storage Platform™, providing the capability to connect or network external, heterogeneous, storage systems, including Hitachi and non-Hitachi systems through standard Fibre Channel port connections, each of which were now virtualized to 1024 virtual ports. This not only provided the ability to connect tens of thousands of heterogeneous host applications, it also provided the ability to network thousands of heterogeneous storage systems together in a common platform of logical, virtual volumes.

Much of the early adoption of the Universal Storage Platform used virtualization as a tool to migrate volumes from old storage systems to new. Today, IT leaders see the cost, performance, protection, and management benefits of deploying tiered storage infrastructures that leverage heterogeneous storage assets and common storage services offered through virtual ports, logical partitioning, provisioning, mirroring, replication, and volume migration. Hitachi also introduced a tiered storage manager that could automate the movement of data across tiers of storage based on policies triggered by time or events.

So, today while other storage virtualization solutions continue to focus on simple volume pooling offering limited value to customers, Services Oriented Storage Solutions from Hitachi Data Systems, enabled by Hitachi storage controller-based virtualization technology, have evolved well beyond basic pooling. Controller-based virtualization allows for common processes, common scripting, common provisioning, common change management, common storage management, common protection, etc. It also allows for a level of heterogeneity that early SANs did not really support. It now provides intelligent tiered storage based upon common management, data mobility, data protection, and provisioning. This capability of virtualized storage systems interacting and working in concert to optimize performance, protection, availability and cost is unique within the industry and can only be accomplished by the purpose-built, multiprocessor storage platform from Hitachi, supported by a large dynamic global cache, and a rich menu of storage and data services.

Extending Services Oriented Storage Solutions to Storage Application Servers

Much of the traditional focus of storage has been around structured data (e.g. databases, transactional data, and data warehouses). However, structured data is becoming a smaller percentage of the overall data storage requirements. Today it comprises only ~20 percent of all data. The other 80 percent is unstructured (such as objects and files) and semi-structured data (such as e-mail), and this unstructured data is growing at close to 10 times the rate of structured data. A recent study from the Enterprise Strategy Group estimates there is close to 2,250PB of unstructured data growing to 20,000PB in 2010. Furthermore, less than 5 percent of this unstructured data is currently being managed by a content management system, and most of the data is dormant after 90 days.

Organizations have come to realize that understanding all data and content is critical to a comprehensive storage strategy. For example, file attributes help classify data, content attributes (metadata) enable additional levels of classification, and the content inside files enables text searching. Recently, there has been a proliferation of niche storage application servers to address these requirements. These storage application servers, including content archive, backup, virtual tape libraries, near-line storage, thin provisioning, and high-performance NAS, all focus on files and content and all have traditionally required their own storage. This requirement created islands of storage, each requiring storage services like replication and migration, and each implementing storage services in their own way. As a result of the need to spend their cycles on storage services, these storage application servers do not scale in performing their primary task, which is to manage and access file and content.

With Services Oriented Storage Solutions, Hitachi centralizes all storage, data, and content services, offloading the workload from the storage application servers and enabling them to concentrate on their application services while also greatly reducing the overall complexity of the storage infrastructure. Recently Hitachi has added the Hitachi Content Archive Platform, which provides indexing, searching, retention, and authentication services specifically for applications requiring active archives. Unlike other content storage systems that require proprietary APIs for access, Hitachi Content Archive Platform services are available through standard protocols, including FCS, NFS, CIFS, HTTP, WebDav, etc. As future storage functions are developed, Hitachi will provide them as services through the Services Oriented Storage Solutions platform. This is analogous to an SOA, where applications can re-use common modules like billing, and not have to rewrite them for every application.

Proof Point for Services Oriented Storage Solutions

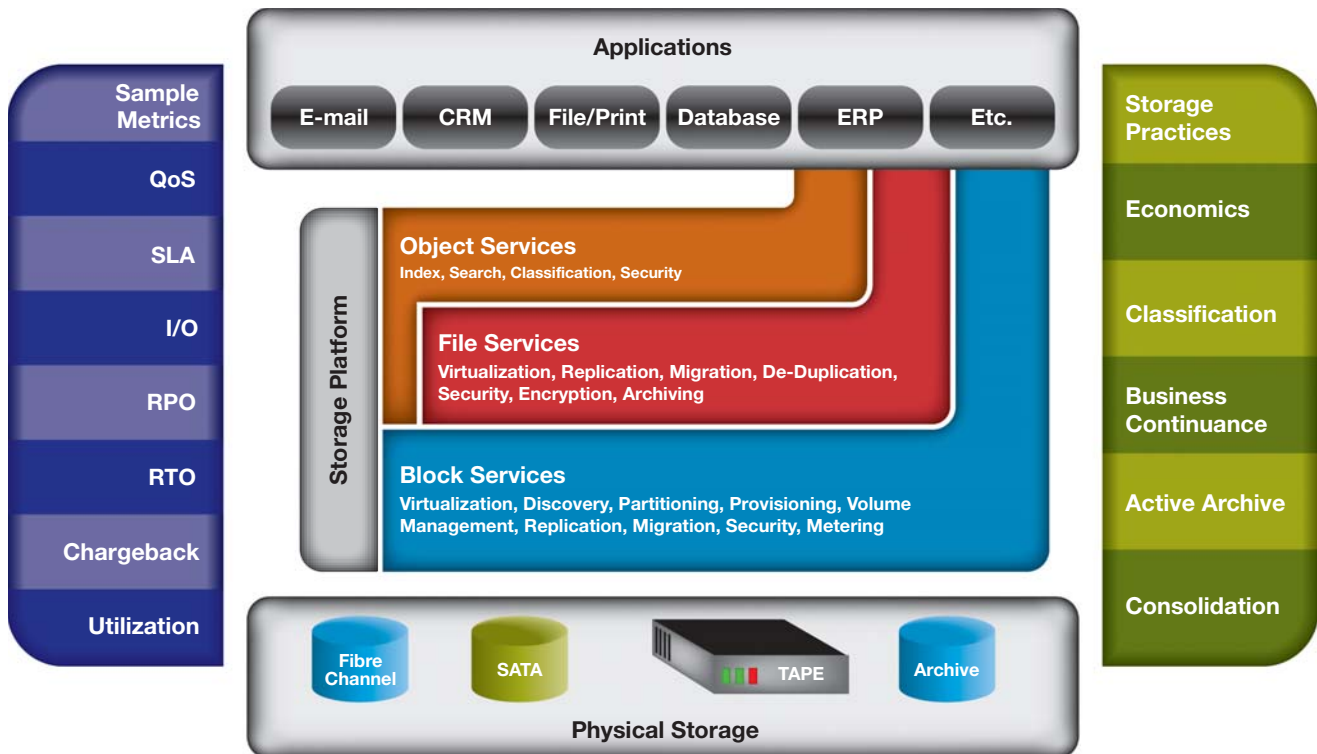
A proof point for this approach is introduction of Hitachi Dynamic Provisioning™ software in the newly announced Hitachi Universal Storage Platform™ V. One of the services that Dynamic Provisioning software provides is thin provisioning, the ability to reserve a maximum capacity for a volume without allocating the actual capacity until it is required. This solves the over-subscription problem of allocated but unused capacity. Normally, this would require a separate standalone storage system like 3ParData, which is specifically designed for this purpose. Since this is provided as a service in the Universal Storage Platform V, this capability will be available for storage applications that attach to it, like NAS, VTL, or content archive. Storage applications do not have to reinvent thin provisioning functionality. They can access this service through standard Fibre Channel or IP protocols. In the future, thin provisioning services will be extended to external storage systems that are virtualized by the Universal Storage Platform V. This ability to extend current and future Services Oriented Storage Solutions to front-end application servers as well as to internal and externally attached storage systems can only be possible on a scalable controller based virtualization platform like the Universal Storage Platform V.

Platform Summary

As we have described throughout this paper, the Services Oriented Storage Solutions platform is a business-centric concept enabling organizations to closely align their storage infrastructure with their business requirements.

While many storage vendors may claim to have business-centric strategies, only Hitachi Data Systems can deliver because Service Oriented Storage Solutions are built upon a dynamic, flexible platform of integrated storage services, enabling customers to optimize storage infrastructure while reducing cost and complexity. The platform is both powerful and simple, as shown in Figure 1.

Figure 1. Services Oriented Storage Solutions comprise an integrated stack of services.



Services Oriented Storage Solutions comprise the following services:

- **Block Services**— which include volume virtualization, discovery, provisioning, partitioning, volume management, replication, migration, security, and metering
- **File Services**— which include file virtualization, replication, migration, security, encryption, and archiving
- **Object Services**— which include content services, such as index, search, classification, and security

These services used individually or collectively deliver Services Oriented Storage Solutions to meet the necessary application storage requirements, based upon metrics including I/O, SLAs, quality of service (QoS), recovery time and recovery point objectives (RTO and RPO), and retention. Most importantly, the unique value of Services Oriented Storage Solutions is the ability to leverage all of these services on a single, integrated storage platform, managed by a common management interface.

Services Oriented Storage Solutions in Action

In addition to hardware and software services components, Services Oriented Storage Solutions offer professional consulting, design, and implementation services to ensure customers maximize their investment in Hitachi Data Systems solutions. The Hitachi Data Systems solution approach is to first understand our customers' key business and IT challenges and then to deploy the appropriate solutions to address their needs. To most effectively help our customers Hitachi Data Systems focuses on the following solution areas:

Storage Management

Today's escalating storage demands challenge you to simultaneously increase capacity, support rapid access to enterprise information, and maintain 24/7 operations—all without additional budget or head count. So, Service Oriented Storage Solutions start with a comprehensive approach to storage management. Storage Management solutions can help you effectively manage your entire storage infrastructure—from the application to the disk spindle—and provide the QoS your business needs. For example:

- **Hitachi Storage Services Manager software** offers services enabling organizations to discover, monitor, visually render, and administer their SAN as a whole, including storage, switches, hosts, and applications from multiple vendors, and provide chargeback based on the resources used.
- **Hitachi Tuning Manager software** includes services to gather detailed information regarding capacity and performance utilization, provide a history for modeling projections for future requirements, and provide alerts based upon high water marks, enabling organizations to automate the process of data migration.

Tiered Storage

A key challenge for organizations is managing storage from multiple vendors and aligning applications with the right storage and then being able to manage data growth, protection, and long-term data retention. Solutions enabled by the **Universal Storage Platform and Hitachi Network Storage Controller™** provide intelligent tiered storage by simplifying infrastructure and reducing costs throughout the enterprise storage network. Advanced virtualization services ensure that these storage systems easily support tiered storage management, enabling organizations to align their application storage requirements with appropriate tiers of storage. For example:

- **Hitachi Tiered Storage Manager, Hitachi Universal Volume Manager, and Hitachi Virtual Partition Manager** offer services that work in unison, allowing you to configure, manage, and move data across the entire storage pool—including volumes on other vendors' storage systems—based upon policies that are triggered by time or events.

Network Attached Storage (NAS)

File serving using NAS is commonplace in today's data center, but organizations are faced with increased challenges around manageability, scalability, and availability. Organizations need NAS solutions that enable them to effectively consolidate and manage all their common file service and high-performance application requirements from a single platform. Hitachi NAS solutions address these challenges:

- The **Hitachi High-performance NAS Platform, powered by BlueArc®**, is an advanced NAS solution comprising file-based services that offer exceptional performance, scalability, and an advanced file virtualization framework for enterprise-class consolidation and high-performance applications. Its policy-based engine for data classification and transparent data migration across storage tiers complements Hitachi Tiered Storage Manager software, allowing organizations to leverage advanced file-based virtualization and block-based virtualization on a single platform.

- Hitachi Data Systems also offers NAS services for the Universal Storage Platform and Network Storage Controller, enabling data center customers to consolidate NAS and SAN onto one footprint.

In addition, Hitachi Data Systems offers NAS services for Hitachi Workgroup Modular Storage and Hitachi Adaptable Modular Storage offering user-scalable file serving at an affordable price.

Business Continuity

Keeping your company operating nonstop is more than just planning for disaster recovery—it's about building a resilient enterprise. That might be as simple as local backup, or as complex as replicating data over long distances. You have to make balanced decisions based on business requirements, including risk exposure, regulatory compliance, best practices, and budgetary constraints. Hitachi Data Systems helps you assess your risks and implement a cost-effective solution that balances cost versus risk mitigation. Hitachi Data Systems offers a comprehensive suite of business continuity solutions including:

- **Virtual Tape Library Solutions by Hitachi Data Systems** enables storage managers to obtain all the benefits of backing up to disk without changing established backup policies or procedures. The integrated suite of storage services creates a “virtual” tape library that appears to the backup application as one or more real tape libraries. Replacing slow tapes with high-speed disk enables customers to increase backup/recovery performance, optimize their storage infrastructure, and achieve their business resilience objectives. Virtual tape libraries can utilize the Universal Storage Platform/Network Storage Controller services for tiered storage and remote vaulting.
- The **Hitachi Data Protection Suite, powered by CommVault®**, is an integrated portfolio of storage services that significantly automates and simplifies key operational and administrative data and storage management tasks. This comprehensive solution, including backup and recovery, disaster recovery, replication, hierarchical storage management (HSM), archiving, and compliance services, is ideal for midsized organizations.
- The **Hitachi Universal Replicator software** delivers simplified asynchronous data replication services across both Hitachi Universal Storage Platform and Network Storage Controller internal and externally attached storage. For organizations that have demanding heterogeneous data replication needs for business continuity or improved IT operations, Universal Replicator delivers the enterprise-class performance associated with storage-system-based replication while providing resilient business continuity, without need for redundant servers or replication appliances.

Active Archiving

Organizations are struggling to comply with regulatory and corporate governance requirements. The risk and exposure to fines and penalties have grown as businesses face increasing scrutiny from government regulators and legal entities. Companies need to be able to prove the authenticity of digital records, provide assurances that they have not been altered, and have the ability to retain records in a protective storage environment over a period of months, years, even decades. While these are key elements for meeting compliance, the capability to easily search for, restore, and retrieve information that may be stored in a number of locations across various media types, such as disk, tape and/or optical, are also critical to ensuring regulatory compliance. Hitachi Data Systems addresses these challenges with:

- The **Hitachi Content Archive Platform**, which allows organizations to satisfy regulatory and corporate governance mandates and provides rapid access to archived data when needed. The solution delivers automated policy management services to ensure content authenticity, retention, and integrity while also supporting robust discovery capabilities through integrated indexing, search, and retrieve services.

Solutions for Smaller Organizations

Most small and mid-sized companies or corporate departments face the same challenges as large organizations (cost, compliance, efficiency, risk), just on a smaller scale; and they need the same kind of storage services as enterprise customers. So, Services Oriented Storage Solutions are designed for organizations of all sizes.

Specifically, the Network Storage Controller, a rack-mounted version of the Universal Storage Platform, is designed specifically for smaller organizations, enabling them to leverage all the benefits of Services Oriented Storage Solutions. For small businesses and departments without a SAN with multiple storage systems, Hitachi Data Systems offers virtualization-in-a-box solutions with modular Adaptable Modular Storage systems. These solutions focus on simplicity and include block and file services, including virtual storage ports for safe application multitenancy, dynamic cache partitioning for QoS, customizable cache block sizes for optimization of database or streaming applications, dynamic multitiering for data lifecycle management, and replication for business continuance. Like all Service Oriented Storage Solutions, Hitachi Data Systems solutions for smaller businesses can be managed using a common management interface.

Services Oriented Storage Solutions Summary

While other vendors may offer individual pieces of the puzzle, no one except Hitachi Data Systems offers an integrated strategy of block, file, and object services providing intelligent tiered storage, common management, data protection, archiving, and performance optimization. In addition, the Hitachi Data Systems strategy is unique in several ways:

- Hitachi Data Systems is the only leading storage company that can scale from very small to very large with a storage approach using a “single pane of glass” management strategy.
- The Hitachi Universal Storage Platform and Network Storage Controller are the only storage platforms in the market that manage heterogeneous storage assets.
- The Hitachi best-in-class virtualization capability is based on an architecture that goes beyond volume pooling to provide a rich set of common storage services for application servers on the front side and heterogeneous storage systems on the back side. This enables storage services to be presented as a unified whole, across block, file, and object data.

Hitachi Data Systems has a singular focus on delivering storage solutions enabling organizations to closely align their storage infrastructure with their business requirements.

By delivering on this promise, Hitachi Data Systems enables organizations to leverage a single universal Services Oriented Storage Solutions platform for delivering storage services and to take an efficient, process-oriented approach to storage, rather than the current, piecemeal, task-oriented approach that leads to needless redundancies, over-subscription of storage resources, management inefficiencies, and compliance exposures.



Appendix A: More Information

For more product information on Hitachi Data Systems refer to:

<http://www.hds.com/products/>

For more information on Hitachi Service Oriented Storage Solutions refer to:

<http://www.hds.com/solutions/index.html>

or

<http://www.hitachisoss.com>



Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA

Contact Information: 1 408 970 1000 info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom

Contact Information: + 44 (0) 1753 618000 info.uk@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., and/or its affiliates in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Universal Storage Platform, Network Storage Controller, Dynamic Provisioning, and Lightning 9900 are trademarks of Hitachi Data Systems Corporation.

Microsoft is a registered trademark of Microsoft Corporation

All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

Hitachi Data Systems sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to http://www.hds.com/products_services/support/warranty.html or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

© Hitachi Data Systems Corporation 2007. All Rights Reserved.

WHP-255-00 LKW May 2007