

# Replication Management

## Overview and Perspectives

A White Paper

*By Christophe Bertrand and Malcolm Muir*

August 2006

## *Executive Summary*

Business continuity has become a priority in the corporate world. IT managers have made replication—local and/or remote—their technology of choice for efficient and effective data protection to ensure business continuity. However, with this choice comes the challenges of ensuring various components and technologies work in unison to deliver the expected service levels and managing an enterprise replication infrastructure that can handle the amount of data being replicated today.

This paper provides an overview of how Hitachi Data Systems is approaching these critical challenges and how its solutions are uniquely positioned (versus the competition) to actively configure, manage, and monitor heterogeneous storage environments and homogeneous Hitachi storage environments, both of which can be virtualized by Hitachi Universal Volume Manager software.

# Contents

<b>Introduction .....</b>	<b>1</b>
<b>Types of Replication Technology.....</b>	<b>1</b>
Local or In-system Replication .....	2
Remote Replication.....	2
Three Data Center Replication .....	3
<b>Hitachi Data Systems Solutions for Managing the Replication Infrastructure .....</b>	<b>3</b>
Monitor, Monitor, Monitor .....	3
Focus on Mainframe .....	6
<b>Top Advantages of Hitachi Data Systems Replication Management Solutions.....</b>	<b>9</b>
Advantages for z/OS Environments .....	9
Advantages for Open Systems Environments .....	10
<b>Conclusion .....</b>	<b>10</b>

# Replication Management

## Overview and Perspectives

A White Paper

By *Christophe Bertrand and Malcolm Muir*

### *Introduction*

Today's business managers recognize data as a business-critical asset. This explains their sudden and intense interest in data protection technologies, in particular those that replicate critical data assets. Thus, business continuity has become a priority in the corporate world, having quickly evolved from disaster recovery planning to a more strategic business governance exercise.

The technology of choice for achieving efficient and effective data protection is replication, be it local or remote. In the enterprise space, and now in the midrange market as well, replication technologies have traditionally been powered by storage systems. So, what are the challenges resulting from this approach? One is that various components and technologies have to work in unison to deliver the expected service levels...and managing an enterprise replication infrastructure is daunting when one considers the amount of data being replicated today, not to mention what future replication requirements will hold.

The ability to successfully endure a crisis or an event that is disruptive to the primary production data center depends to a large extent on the organization's ability to protect data in a way that limits the impact of disruption to the business. This impact is usually quantified in terms of recovery point objective (RPO) and recovery time objective (RTO): *RPO* is defined as being the point in time to which data must be restored to successfully resume processing (often thought of as the time between last backup and when the outage occurs); *RTO* is defined as the timeframe in which business functions or applications must be restored, and it is quantified by the time after the disaster is declared through to the time it takes to perform the recovery tasks.

This paper provides an overview of how Hitachi Data Systems is approaching this critical data protection challenge and how its solutions are uniquely positioned (versus the competition) to actively configure, manage, and monitor heterogeneous storage environments and homogeneous Hitachi storage environments, both of which can be virtualized by Hitachi Universal Volume Manager software.

### *Types of Replication Technology*

Common questions asked by those considering an investment in replication technology are:

- :: "What mode of replication technology should I use?"
- :: "Should it be synchronous or asynchronous?"
- :: "What are the tradeoffs?"
- :: "Which one provides the best data protection option for me?"

## Local or In-system Replication

A first line of protection can be local or in-system replication, a process that copies volumes of data internal to the storage system. This allows point-in-time copies to be leveraged for a variety of needs, well beyond those of data recovery, addressing such business challenges as backups, change/control, maintenance, use by other applications for analytical purposes, and more. More importantly, using a local full-volume copy on a cyclical basis also protects users from data corruption, allowing them to revert to an uncorrupted version of their data. This is a critical function for mission-critical applications, such as e-mail, and high-transaction applications, such as online transaction processing (OLTP).

## Remote Replication

If site-wide events occur, however, data assets are exposed to destruction. To mitigate this risk, replicating data over a certain distance is the answer. Two remote replication modes traditionally used are synchronous remote replication and asynchronous remote replication.

### Synchronous Replication

With synchronous replication, the storage system “intercepts” the data being written by applications. Before these updates are written to the primary storage system they are sent over a communication link to a remote storage system, after which the write is then acknowledged as being complete to the source storage system. This same write is then written to the primary storage system and the application is then sent an acknowledgement of the write having been completed. Since the data is applied to the remote system first, the replicated write is guaranteed. The storage system handles the overall process with built-in mechanisms to guarantee that transactions are indeed applied to both systems.

As data is sent to the remote system first, a delay, referred to as latency, is introduced because the process requires that the update first be copied to the secondary site. The duration of this latency is affected by the volume of transactions, network factors, and the distance between the primary and remote sites.

Typically, this means that synchronous replication will be used in metropolitan areas where data centers are located no more than 60 miles apart, although longer distances are possible. A major benefit is the ability to keep revenue-producing applications online, virtually without interruption. This is made possible with clustering technology and synchronous remote replication, and it permits applications to fail over in real time and have support personnel locally available to staff at the remote center. One of the major drawbacks of this approach is that it exposes an organization to a regional outage or disaster.

### Asynchronous Replication

Asynchronous replication technology, on the other hand, permits application updates to be written to disk immediately on the primary system, after which these same updates are then sent to the remote site—hence, the asynchronous nature of the process. Key challenges that may stem from this approach include ensuring that the updates written to the secondary site are in the right sequence (application writes can be dependent upon each other), and determining how far behind the remote site is from the real-time transactions occurring on the primary site. Administrators will also consider the RPO.

Proven asynchronous replication technology will include metadata with the data being replicated in order to guarantee its integrity and, therefore, guarantee that data on the remote site is consistent when the application is restarted or, at the very least, recoverable. Without this proper sequence of writes, applications can become corrupt and unusable.

One of the major benefits of asynchronous replication is the ability to replicate over long distances. Therefore, organizations can protect their data “out of region,” sometimes thousands of miles away, at a secondary or remote site. Furthermore, with modern and proven asynchronous replication technology it is possible to achieve an RPO of a few seconds (being only a few seconds behind). One assumption is that one can actually measure the RPO accurately and with granularity.

## Three Data Center Replication

When synchronous and asynchronous replication is combined into what is referred to as “three data center” replication solutions, maximum protection and flexibility in recovery is achieved. Recent developments in storage-based replication technology make this approach possible. A critical dimension of the solution will prove to be how successfully one can control the replication processes of what might be thousands of volumes, across multiple sites, over short and long distances.

While these technologies can certainly do the job and address the protection of data issue, monitoring the replication environment and actively managing it has become critical. A poorly managed replication infrastructure will fail to fulfill its mission. As the old saying goes, “The proof is in the pudding.” While there is really no culinary dimension to IT-focused business continuity, the saying exemplifies that unless and until one can prove that data can be recovered in a consistent manner within a certain amount of time following a variety of events or disruptions...it will be difficult to measure a successful replication implementation.

For example, in three data center replication solutions, the replication environment has become more and more complex to manage in a way that meets demanding RPO and RTO requirements. In the production enterprise data center, the replication topology to manage would consist of hundreds or thousands of replicated pairs and different types of replication technologies (local or remote, synchronous or asynchronous). And, more challengingly, all of this would be deployed across multiple locations with distance. Simplifying and optimizing the replication management function has actually become as critical as providing the actual replication services.

## *Hitachi Data Systems Solutions for Managing the Replication Infrastructure*

Hitachi Data Systems offers a number of replication solutions for both local mirroring (full volume or snapshots) as well as distance replication. These software solutions run within the storage system microcode and are managed through an innovative replication management framework integrated with Hitachi Storage Management and HiCommand® Suite software. These tools not only provide high-level monitoring and reporting capabilities, but also support advanced remote management functions for select platforms.

## Monitor, Monitor, Monitor

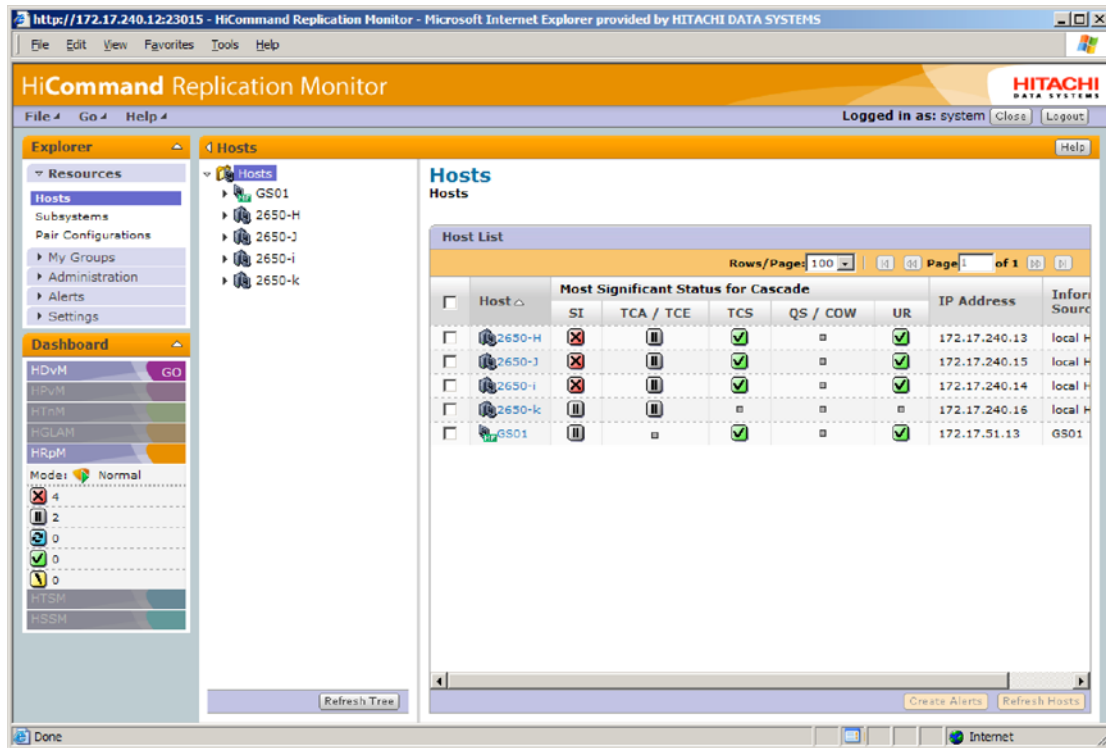
This seemingly simplistic mantra is actually not so simple, and it is key to storage administrators getting a better night’s sleep! There is no doubt that without a clear understanding of the health and status of complex replication infrastructures, companies cannot be comfortable that their number one asset—data—is being protected.

With a single, easy-to-use, centralized management console for monitoring and visualizing volume replication configurations and status information, HiCommand Replication Monitor software simplifies

administration of the entire suite of Hitachi replication products for open systems and mainframe environments.

Replication Monitor software helps integrate with HiCommand Device Manager software (where replicas are defined) and also with the replication management functions of Hitachi storage systems. This collaboration allows storage administrators to get a visual reference for applications being replicated (see Figures 1 and 2), as well as a point-in-time status indicator of replication pairings.

**Figure 1. HiCommand Replication Monitor Software's Centralized Management Console Screen—Enterprise Environment**



*HiCommand Replication Monitor software simplifies administration of the entire suite of Hitachi replication products for open systems and mainframe environments.*

### The Key Role of Centralized Management

HiCommand Replication Monitor software now provides a single tool to review the enterprise-wide replication environment, regardless of the operating system:

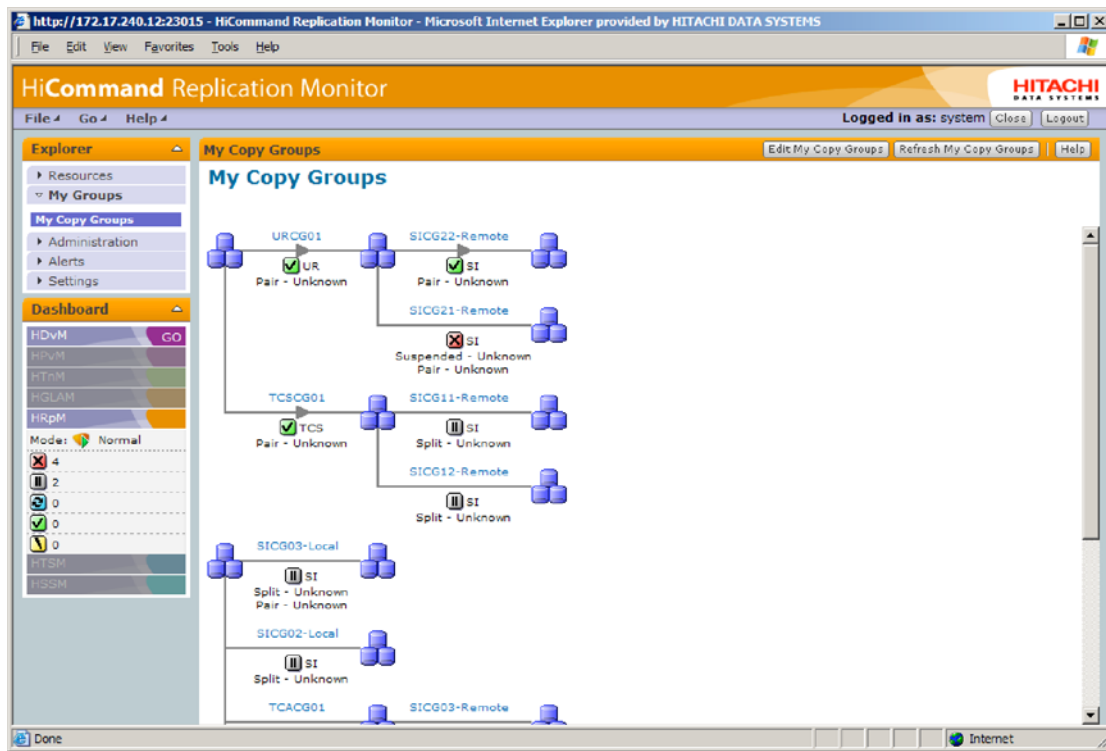
- :: Easily oversee open and mainframe replicated environments
- :: Quickly determine all the associated pairs/groups of selected resources
- :: Quickly identify potential problems
- :: Easily employ SNMP and e-mail alert mechanisms
- :: Centrally monitor utilization of replication sidefiles and/or journals [Monitoring these replication components is critical to the health of distance replication (asynchronous) solutions.]

## Focusing on Usability and Flexibility

Replication Monitor software now provides an enhanced topological-like view of user-selected copy groups:

- :: The “My Copy Groups” feature allows administrators to create their own topological view of the replicated environments. In addition, users obtain an immediate visual summary of the various complex replication environments.
- :: Replication Monitor software now supports a common user interface, fully integrated with the HiCommand storage management framework.

**Figure 2. HiCommand Replication Monitor Software’s Centralized Management Console Screen—“My Copy Groups”**



The “My Copy Groups” feature allows administrators to create their own topological view of the replicated environments.

## Measuring Data Currency and Recovery Point Objectives

Given today’s complex replication environments, it has become critical that customers monitor how far “behind” the data is at the remote site. Replication Monitor software now permits administrators to monitor this lag with a very high level of granularity:

- :: Provides numbers for both open and mainframe groups
- :: Makes historical data available for download; eases planning, reporting, and managing the currency of replicated data
- :: Sends e-mail and/or SNMP trap alerts when an RPO threshold is bypassed

Replication Monitor software provides centralized management of Hitachi Remote Replication and Hitachi In-System Replication software. This ability, in turn, decreases management complexity while increasing staff productivity and provides greater control than was previously available. This results in the improvement of service levels by reducing planned and unplanned downtime of customer-facing applications.

For business managers, Replication Monitor software serves as an RPO watchdog, ensuring efficient and effective business continuity for data protection management. Its capacity to monitor data currency and recovery points for both open and mainframe environments enables critical advantages. In addition, leveraging advanced replication system status alerting and reporting provides storage administrators with early warnings of potentially larger problems—thereby allowing them to be proactive versus reactive. Problem-avoidance is built-in and a reduction in outages can also be achieved through attributes such as the powerful auto-discovery feature. This feature enables enterprise-wide storage configuration and replication objects to be identified and monitored. It should also be noted that Replication Monitor software assists in the reduction of reconciliation times of pair management should outages occur, which from time to time they do...unfortunately!

## Focus on Mainframe

In the mainframe world, traditional approaches to setting up, managing, and recovering from outages can be labor-intensive, as the storage administrator has had to rely mostly on JCL input statements and/or the TSO command line. HiCommand Business Continuity Manager software for IBM® z/OS® provides centralized, simplified, and automated enterprise-wide management of Hitachi Remote Replication and In-System Replication software for mainframe environments.

Through a single, consistent, interface based on familiar TSO/ISPF dialog screens (see Figure 3), Business Continuity Manager software automates Remote Replication and In-System Replication software operations. In addition, monitoring for this heterogeneous operating system enterprise is provided through the Replication Monitor software, which communicates with Business Continuity Manager software to enable a “single pane of glass” monitoring tool.

**Figure 3. Business Continuity Manager Software's Interface/Dialog Screen**

```
                                Add Pair Extent
                                2006/06/21 17:05:48
Hitachi Business Continuity Manager 5.0 -----
Command ==>
Select a device number or volser with "/". Specify a range of device
numbers at the Secondary Site. Then press Enter.
Copy Group ID . . : CFGA1

                                Primary Site: PDAD          Secondary Site: SDAD
                                240 devices specified

Devn ( / ) | Start Device   : x 5000      Start Device: x C000
            | End Device     : x 50EF      End Device   : x C0EF

Volser (   ) | Start Volser   :
            | End Volser     :
            | SMS STR GRP Name:

F1=Help   F3=Exit   F12=Cancel
```

*Through a single, consistent interface based on familiar TSO/ISPF dialog screens, Business Continuity Manager software automates Hitachi Remote Replication and In-System Replication software operations.*

Access to critical system performance metrics and thresholds allows proactive problem avoidance and optimum performance to ensure that service-level objectives are met or exceeded. Business Continuity Manager software's auto-discovery capability eliminates hours of tedious input and costly human error when configuring and protecting complex, mission-critical applications and data.

Business Continuity Manager software is a feature-rich tool that helps administrators define copy groups (mostly application-based) that contain multiple replication objects with similar attributes so that they can be managed with a single command. With the number of application volumes that needs to be managed for replication, this control is imperative and delivers administrative scalability. This auto-discovery feature also eliminates manual errors and streamlines management for all replication objects.

**Figure 4. Business Continuity Manager Software—Manage Copy Groups Screen**

```

                                Add Pair Extent
                                2006/06/21 17:05:48
Hitachi Business Continuity Manager 5.0 -----
Command ==>
Select a device number or volser with "/". Specify a range of device
numbers at the Secondary Site. Then press Enter.
Copy Group ID . . : CFGA1

                                Primary Site: PDAD           Secondary Site: SDAD
                                240 devices specified

Devn ( / ) | Start Device   : x 5000           Start Device: x C000
            | End Device     : x 50EF           End Device   : x C0EF

Volser (   ) | Start Volser   :
            | End Volser     :
            | SMS STR GRP Name:

F1=Help   F3=Exit   F12=Cancel
```

*Business Continuity Manager software allows administrators to define copy groups (mostly application-based) that contain multiple replication objects with similar attributes so that they can be managed with a single command.*

Among many other functions, Business Continuity Manager software permits users to view the status of all enterprise-wide replication objects in real time and accesses key replication metrics with built-in performance monitoring. It also provides automatic notification of key events completion, such as pair state transitions, timeout thresholds, and other system events. It should be noted that Business Continuity Manager software employs technology developed by Hitachi to deliver nondisruptive, periodic, point-in-time remote data copies across any number of heterogeneous storage systems and over any distance.

### Business Continuity in Action

One of the key exercises to ensure business continuity success is periodic testing of recovery plans. Executing the test plan, which includes running the processes by the responsible people *and* making sure one can actually recover the data and activate the applications, is paramount to success. Business Continuity Manager software enables administrators to test and implement frequent, nondisruptive disaster recovery testing with current and accurate data, in the form of an ISPF/PDF user dialog. And when a disruptive event occurs, it is critical to have already identified and automated these processes. Business Continuity Manager software automates complex disaster recovery and planned outage functions. This results in dramatically reduced recovery times. Note: Business Continuity Manager software is also used to enable three data center disaster recovery configurations.

### Simplifying for Better Results

Being proactive is absolutely essential, and it requires the ability to rely on built-in capabilities for monitoring. This forms the foundation for active management of critical performance metrics and thresholds that support proactive problem avoidance. As mentioned earlier, Business Continuity Manager software delivers auto-discovery of enterprise-wide storage configurations and replication objects, which helps in eliminating tedious, error-prone data entry that can cause outages. Finally, because IT environments are all unique in how they support their business, Business Continuity

Manager software also simplifies the development of additional automation with built-in automation and extensibility features that can be leveraged using the familiar REXX language. As a result, this capability ends up enabling the protection of more data with less time and effort (also known as “administrative scalability”).

### Focusing on Productivity and Process Improvements

Business Continuity Manager software clearly decreases management complexity while increasing staff productivity and provides greater control than previously available. Furthermore, it allows users to significantly accelerate deployment of Remote Replication and In-System Replication software solutions through a single, consistent user interface based on the familiar TSO/ISPF full-screen panels. This also shortens time to deploy new business applications. Overall, it supports the business infrastructure and mission by improving service levels through reduced planned and unplanned downtime of customer-facing applications.

## Top Advantages of Hitachi Data Systems Replication Management Solutions

The exponential growth of data and the expanding obligations to protect it have heightened the importance of providing automated and scalable storage management solutions. In response to new legal requirements and governmental regulations as well as the increasing business value of information, the tasks of replicating data and managing replicated data have been pushed to the forefront of business considerations and decisions. It is for this reason that Hitachi Data Systems is placing a high level of importance on providing automated and scalable replication management solutions.

### Advantages for z/OS Environments

- :: The Business Continuity Manager copy group containers feature enables storage administrators to submit replication commands to multiple consistency groups at the same time. This feature reduces the number of lines of code required in TSO command procedures and REXX Execs.
- :: Business Continuity Manager software is far more efficient than comparative solutions, executing a single management command per consistency group, versus a command per pair, as is required by comparative solutions. This results in a much higher level of efficiency and less TSO/REXX Exec scripting requirements.
- :: Business Continuity Manager command devices support communication, command and control, to all related storage systems from a single Business Continuity Manager image, which also includes remote storage systems.
- :: Business Continuity Manager software is optimized for performance. With the amount of data and number of replicas that require management today, optimized performance is mandatory. Comparatively speaking, a 20-fold improvement in both elapsed and CPU times can be experienced when executing mass operations—in some cases thousands of volumes are involved.
- :: Business Continuity Manager software affords a storage administrator the flexibility to use either Business Continuity Manager software or the graphical user interface (GUI) presented by Replication Monitor software. For mixed environments (mainframe and open systems) Replication Monitor software is especially useful.

- :: Often, IT must adhere to and account for stringent service level objectives. Business Continuity Manager software provides a metric that measures application RPOs. Although comparative solutions also provide this metric, they cannot provide as accurate a measurement, as Replication Monitor software benefits from the underlying architectural features found only in asynchronous solutions by Hitachi, such as I/O time stamping, for example.

## Advantages for Open Systems Environments

- :: Today, as open systems applications take on a mission-critical role in business operations, IT administration staff members are being held accountable to adhere to published service level objectives. As with the mainframe environment, Replication Monitor software provides a measurement for RPO.
- :: Replication Monitor software is a storage management module in the HiCommand Suite of software and integrates with the HiCommand framework. Replication Monitor software obtains pertinent device management metadata from the HiCommand Device Manager software database—and, in the case of mainframe environments, obtains replica information from Business Continuity Manager software, consolidating both to a centralized management console. No other comparative solution is moving in such an all-inclusive and integrated direction.
- :: Comparatively speaking, the usability offered by the Replication Monitor software user interface is second to none, with the same look and feel and alert notifications presented for both mainframe and open systems replication environments.
- :: The most recent release of Replication Monitor software provides support for Device Manager logical groups. In most cases, Device Manager logical groups represent applications. This gives a whole new administrative scalability capability to Replication Monitor software, as views and actions can be taken on behalf of the business application. This represents yet another advancement supporting Application Optimized Storage™ solutions from Hitachi Data Systems.

## Conclusion

Hitachi pioneered the development of asynchronous remote data replication and continues to improve this capability. The significant advantages afforded by the combination of Replication Monitor software and Business Continuity Manager software exemplify the direction in which replication management is going. The HiCommand Suite includes these software modules as well as replication management software products that are key components of Application Optimized Storage solutions for business continuity.

By raising the bar yet one more time, and providing a “single pane of glass” for monitoring the replication infrastructure, Hitachi is setting the stage for the next generation of replication management tools. This is the key to improved data asset management and protection, and a better night’s sleep for many IT professionals!

 **Hitachi Data Systems Corporation****Corporate Headquarters**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
**www.hds.com**  
**info@hds.com**

**Asia Pacific and Americas**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
**info@hds.com**

**Europe Headquarters**

Sefton Park  
Stoke Poges  
Buckinghamshire SL2 4HD  
United Kingdom  
Phone: + 44 (0)1753 618000  
**info.eu@hds.com**

No part of this document may be reproduced or transmitted without written approval from Hitachi Data Systems, Inc.

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd. HiCommand is a registered trademark of Hitachi, Ltd.

Application Optimized Storage is a trademark of Hitachi Data Systems Corporation.

All other product and company names are, or may be, trademarks or service marks 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/license.html](http://www.hds.com/products_services/support/license.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.

©2006, Hitachi Data Systems Corporation. All Rights Reserved.

WHP-226-00 August 2006