

# Enhancing IBM<sup>®</sup> Tivoli<sup>®</sup> Storage Manager with Virtual Tape Library Solutions by Hitachi Data Systems

An In-depth Look at Supporting Benefits of Diligent<sup>™</sup> ProtecTIER<sup>™</sup> Technology

Application Brief

May 2007



## Executive Summary

With disk prices falling and the need to recover data quickly increasing, more and more organizations have started adopting disk as a viable alternative to slower, more resource-intensive magnetic tape in their backup and recovery environments. Supporting these disk alternatives is IBM® Tivoli® Storage Manager, one of the leading enterprise backup and recovery applications available on the market today.

Virtual Tape Library Solutions by Hitachi Data Systems provide the ideal disk storage platform for Tivoli Storage Manager. With its advanced HyperFactor™ data deduplicating technology, ProtecTIER™ from Diligent™ Technologies Corporation, enables Virtual Tape Library Solutions to greatly improve the speed and reliability of backup and recovery operations while significantly reducing the total cost of ownership of a disk-based backup and recovery environment.



# Contents

Introduction .....	1
Customer Challenges .....	1
Disk May Be the Solution .....	1
Increasing Backup Performance .....	2
More Importantly, Increasing Restore Speed .....	2
Reducing Costs .....	3
Improving Reliability.....	4
Assuring Data Security.....	4
Enabling Advanced Data Protection Scenarios.....	5
Seamless Integration with TSM.....	6
Moving Data to Tape .....	7
Enhancing Progressive Incremental Backups .....	8
Improving Collocation .....	8
Improving Reclamation.....	9
Conclusion .....	9



# Enhancing IBM® Tivoli® Storage Manager with Virtual Tape Library Solutions by Hitachi Data Systems

An In-depth Look at Supporting Benefits of Diligent™ ProtecTIER™ Technology

Application Brief

## Introduction

Tivoli® Storage Manager (TSM) from IBM®, is one of the leading enterprise backup and recovery applications available on the market today. TSM provides features and capabilities that help thousands of backup administrators protect the mission-critical data that is the lifeblood of their organizations. With the rapid data growth that organizations are experiencing, even the best backup application needs a little help.

Riding to the rescue is Hitachi Data Systems, an IBM partner. Virtual Tape Library Solutions by Hitachi Data Systems employ powerful ProtecTIER™ technology from Diligent™ Technologies Corporation to enhance and improve TSM's ability to protect mission-critical data without disrupting existing operations.

## Customer Challenges


According to a 2005 study conducted by the independent storage industry analyst firm, Enterprise Storage Group, organizations are continuing to struggle with their tape-based backup and recovery environments. The biggest issues IT organizations are facing include:

- Backups that take too long
- Recoveries that take even longer
- Inability to validate data recoverability
- Resources required to manage tape media

## Disk May Be the Solution

With disk prices falling, organizations have started adopting disk as a viable alternative medium for storing backup data. Backup and recovery from disk can solve all of these challenges. An effective disk-based backup and recovery environment requires the backup and recovery application to know how to take advantage of the random access capabilities of disk.

Although TSM has supported disk storage pools for many years, there are many advantages to using a virtual tape library in place of, or as a supplement to, your disk pools. For example, you don't need to assign a dedicated disk pool for every server. With the ProtecTIER enterprise-class data protection platform you can either share a Virtual Tape Library Solution among many servers or you can create several virtual tape libraries, one or more for each server. This means you are not required to dedicate storage to every server, because the Virtual Tape Library Solution that uses ProtecTIER technology can assign capacity as required.



More importantly, ProtecTIER's unique HyperFactor™ data de-duplicating technology provides up to 25 to 1 reduction of storage space required, delivering the capacity to retain months of backup images at a total cost of ownership (TCO) less than that of actual tape.

ProtecTIER technology was designed from the ground up to work seamlessly with TSM to offer remarkable improvements in speed and reliability and to greatly reduce the TCO of a disk-based backup and recovery environment.

## Increasing Backup Performance

In a tape-based backup and recovery environment, backup administrators need to carefully tune TSM to stream just the right amount of data to a tape drive. Streaming too little data can result in “shoe shining,” an effect of starting, stopping, and repositioning the tape that results in poor drive performance. Multiplexing and parallel sessions are used to make sure enough data is being sent to maximize tape drive utilization, compensating for slow clients, slow networks, or small backup jobs.

ProtecTIER frees backup administrators from shuffling backup jobs and tuning tape libraries and operations for optimal performance. Regardless of the amount of data being streamed, virtual tape libraries back up data at the speed of the target disk.

In addition, although virtual tape libraries accept and process all tape operation commands, they are not hit with the delays associated with allocating tape cartridges, robotic movements, media mounting, media positioning, or media eject operations. In a tape-based environment, these delays can last from a few seconds to a few minutes per tape.

Increased backup performance allows IT organizations to complete backups within allowed backup windows and meet service level requirements of mission-critical applications. With ProtecTIER, optimizing backup performance is no longer achieved at the expense of recovery performance.

## More Importantly, Increasing Restore Speed


As discussed previously, backup administrators use multiplexing to ensure a constant stream of data flows to each tape drive in order to reach optimal drive performance. Multiplexing is often the only way to back up all data in a given backup window with the tape drive resources provided.

Taking advantage of TSM's multiplexing capabilities and parallel sessions increases backup performance but has a devastating impact, slowing down restore performance significantly. This is primarily due to the extra time required to read through backup images and skipping over unwanted data from other backup jobs combined into that image.

In addition, TSM's progressive incremental backup methodology results in data being dispersed over a large number of tapes. Depending upon how long backups have been running using TSM, the data needed to restore a file system could span 30 or more different tapes. The extensive time spent moving, loading, mounting, positioning, and unmounting cartridges, multiplied by dozens of tapes can result in extremely long restore times.

ProtecTIER eliminates the delays associated with moving, loading, and mounting of cartridges and the seek times inherent to physical tape.

By recovering data from ProtecTIER at the speed of disk, TSM can deliver an order of magnitude improvement in recovery performance and, more importantly, reduce downtime and its associated costs.



One TSM customer who turned to Hitachi Data Systems for help was experiencing some significant performance issues with their tape-based backup environment. Recovering data from tape had become increasingly time-consuming. For example, prior to implementing the Virtual Tape Library Solution with ProtecTIER they had the occasion to recover 1 million files stored on 10GB of tape and it took them more than 17 hours to complete. After installing the Virtual Tape Library Solution with ProtecTIER and integrating into the TSM environment, they needed to recover another 1.6 million files on 34GB of disk. This task took less than 1 hour and 38 minutes to complete, demonstrating the remarkable improvement in performance that can be gained by integrating a Virtual Tape Library Solution with ProtecTIER into a TSM environment.

In addition, according to Strategic Research Corporation, an independent analyst firm, the vast majority of all restores are of single files, not full file system recoveries. Our performance tests show that restoring a single file from ProtecTIER is hundreds of times faster than restoring from even the fastest tape drive. This is due to the ability to randomly access data on disk and the elimination of mounting, positioning, and unmounting delays associated with tape-based backups. ProtecTIER enables single file restore operations to be completed in the time it normally takes for a physical tape to be mounted.

## Reducing Costs

It is obvious that backing up and recovering from disk is much faster than doing so from tape. Yet, even knowledgeable IT professionals have difficulty grasping how introducing disk into a backup environment can reduce costs. And, it is true that when using TSM's disk-to-disk (D2D) capabilities with standard disk-based systems it is too expensive to store more than a small subset of critical backup data. Most enterprises need to recover data for weeks or months, which means many versions of files or databases must be saved on a daily basis. To save multiple versions in a standard disk storage pool becomes prohibitively expensive.

However, ProtecTIER's unique, advanced, and unmatched HyperFactor data de-duplication capabilities reduce physical disk storage requirements by up to 25 times or more and enable TSM to provide quick recovery of data from disk at a cost less than that of tape.

Most TCO studies only compare the cost of automated tape libraries and their associated media to the cost of large disk storage systems with virtual tape library servers and software. These hard costs are only part of the total equation. These TCO studies are weak on operational cost and resort to making generalized assumptions. Strategic Research Corporation recently completed a study that includes hard costs as well as the cost to administer failed backups (that first-time backup error rate problem that tape struggles with), labor cost of recovery, vaulting, tape management, etc. The research data shows an average 30 percent reduction in the operational costs of administering backup and recovery across the 10TB to 100TB environment size just by moving from tape to a disk-based virtual tape library.

Combining the data protection capabilities of TSM and Virtual Tape Library Solutions with ProtecTIER reduces the time and cost associated with managing backup and recovery processes, and eliminates cumbersome tape media management and maintenance activities like tape head cleaning and tape retensioning.

There is no doubt that need to maintain data on physical tape media will continue for the foreseeable future. However, supplementing TSM with ProtecTIER technology can enable IT managers to consolidate tape library resources, lowering hardware replacement costs, implementation and training costs, remote vaulting services costs, and significant annual drive and library maintenance costs.

## Improving Reliability

If you're looking for the best way to increase the reliability of your TSM environment, begin by reducing the cause of most failed backup jobs, tape.

Tapes, tape drives, and the robotic tape libraries are mechanical devices and are prone to failures, including:

- Tape drive failure
- Media errors
- Tape getting stuck in a drive
- Can't read tape due to aging or normal wear and tear
- Tape not loaded
- Wrong tapes loaded
- Tape demagnetized
- Robotic arm jam

The reliability of tape is a big issue and frequently causes backup jobs to fail, requiring manual intervention and time spent restarting jobs to maintain the necessary level of data protection.

Another TSM customer who looked to Hitachi Data Systems for assistance was experiencing some significant issues with their tape-based backup environment. Their backups were failing with great regularity (20 percent to 30 percent of all jobs), to the point where they had a full-time employee who did nothing but identify and repair the failed backups. They were unable to complete their backups on a daily basis and believed tape reliability issues caused most of the failures.


And as their IT manager said, "Our shop was exposed every day in some environments." They couldn't protect their mission-critical data.

Since implementing ProtecTIER, all tape related failures were eliminated, resulting in nearly 100 percent of all backup jobs completing successfully. Their critical data is now fully protected and the employee who struggled with failed backup operations has moved on to provide more valuable services to the company.

## Assuring Data Security

Human intervention accounts for the majority of security risks associated with physical tapes. These risks include tape loss, theft, or damage during transportation to offsite vaulting facility. Based on the size of data and the multiplexing and incremental techniques used, a single restore operation could require retrieving data from dozens of tapes. Losing just one of those tapes may make the entire restore operation fail, requiring the backup administrator to rely on older backup jobs. This may result in the loss of a significant amount of data and greatly compromise the overall Recovery Point Objective (RPO).

ProtecTIER can be used in conjunction with TSM to eliminate the risks of errors associated with tape handling and media management. Although robotic tape libraries automate some operations, manual handling of tape cartridges is still required and risks data loss due to mislabeling, misfiling, or mishandling. With a Virtual Tape Library Solution, administrators no longer need to worry about these risks.



ProtecTIER can also be used to eliminate human intervention and its risks associated with the transportation of physical tape to an offsite storage facility. Using a remote replication technology, such as Hitachi TrueCopy<sup>®</sup> Synchronous, Hitachi TrueCopy Asynchronous, or Hitachi Universal Replicator software, ProtecTIER's virtual tape cartridges can be electronically vaulted over the wide area network (WAN) to a remote storage facility or disaster recovery site, providing complete protection throughout the entire data protection cycle, from backup to vaulting. Since ProtecTIER greatly reduces the amount of physical storage needed by a factor of up to 25 times or more, the major cost of the WAN bandwidth required is significantly decreased.

TSM can be configured to automatically vault virtual cartridges immediately after creation. By automating this process, data is immediately protected without human intervention and IT staff resources are free to focus on more strategic tasks.

Virtual Tape Library Solutions also protect data from unrecoverable losses by striping data across multiple disks. Physical tapes have a limited life expectancy. When a tape cartridge wears out, is damaged, or fails, the data on it is lost. On the other hand, a RAID disk array has an almost unlimited life expectancy. If a disk drive fails, it can be replaced without losing any of the data stored on the RAID array. TSM policies for tape life can be increased because the tape life is now based on virtual tapes not physical tapes.

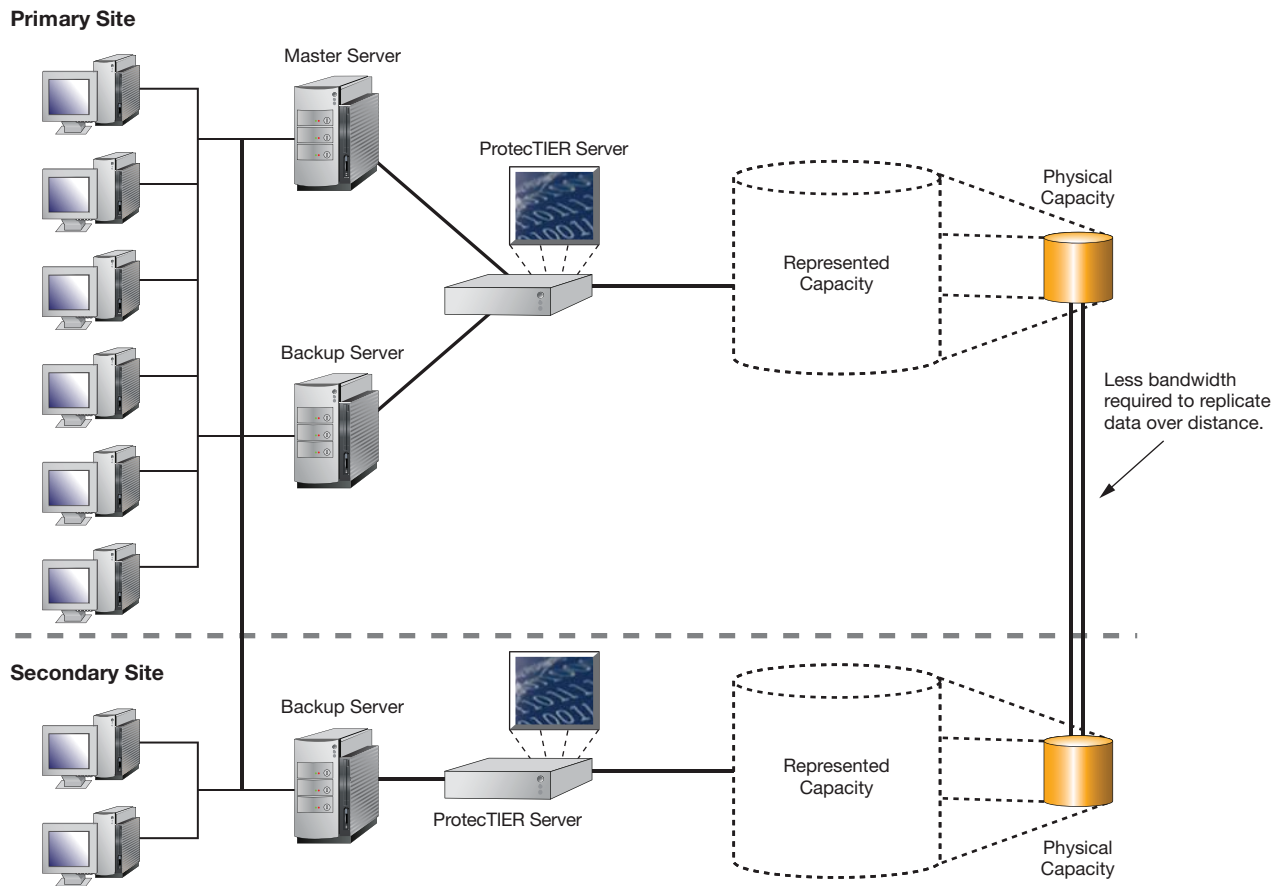
## Enabling Advanced Data Protection Scenarios

In most TSM environments, backup data is written to tape and physically transported to offsite storage or a hot-site location for disaster recovery protection. This antiquated, expensive, and manually intensive operation is required because electronically transporting the huge amount of backup data generated daily is cost-prohibitive.

With ProtecTIER's advanced HyperFactor data deduplicating capabilities, backing up and recovering data to and from remote locations can be done quickly and can become an affordable reality. By significantly shrinking the amount of data sent by as much as 25 times or more, ProtecTIER technology significantly reduces the bandwidth required to electronically vault data to remote locations by transmitting only new and unique data segments across a WAN. This enables IT organizations to replicate selected mission-critical backup data from a primary location over a WAN to a secure offsite location. The electronic vaulting enabled by the massive data reduction provided by ProtecTIER technology significantly reduces the bandwidth costs, which is the major cost component of wide area data transfers. ProtecTIER makes using networks to move backup data offsite for disaster recovery protection both economically and operationally feasible.

In addition to the site-to-site recovery model shown in Figure 1, businesses can implement automatic electronic replication in several ways: one-to-one, one-to-many, or many-to-one, depending on their environment and desired business continuity goals.

**Figure 1. Site-to-site Recovery**



*ProtecTIER technology, in a Virtual Tape Library Solution, makes using networks to move backup data offsite for disaster recovery protection both economically and operationally feasible.*

## Seamless Integration with TSM

By emulating an actual tape library, ProtecTIER enables TSM to function identically with virtual tape as it does with physical tape, except that it does it much faster and much more reliably.

TSM administrators can enjoy the benefits of more reliable backups and faster restores, without changing any of the existing policies and procedures they have worked so hard to implement.


Using the standard administration graphical user interface (GUI), a TSM administrator can quickly change the target of a backup job from an existing automated tape library to a virtual tape library (VTL). A few clicks are all that is needed to dramatically improve the performance and reliability of backup operations.

In contrast, setting up and maintaining file systems to use as the targets for native TSM D2D backups can be time-consuming and can result in wasted storage space and increased costs. Setting up and maintaining disk pools takes a lot of time because each backup job requires its own file system to be created and maintained. This may be acceptable for a small number of backup jobs, but the administrative burden associated with hundreds or thousands of backup policies can quickly become a full-time job. Additionally, the backup administrator often does not have the expertise or authorization to set up file systems and must rely on assistance from a storage administrator.

Using file systems as the targets of D2D backups can also result in a large amount of unused and wasted disk space. When setting up a disk target the storage administrator must choose the amount of storage allocated. If sized too small, the backup jobs may fail. If sized too large, the unused space is wasted. In most cases, the repository is sized considerably too large, resulting in a significant amount of storage being wasted.

ProtecTIER, on the other hand, efficiently uses storage space and maximizes utilization, resulting in no wasted storage. (See Figure 2.) In addition, all of the storage management tasks required are hidden by the virtual tape engine and are automatically and transparently accomplished. After installation, special storage management expertise is no longer required for ongoing maintenance purposes.

**Figure 2. Virtual Tape Library Solution with ProtecTIER—No Wasted Storage**




*Implementing ProtecTIER in a Virtual Tape Library Solution is nondisruptive and can immediately improve backup and recovery reliability and performance.*

## Moving Data to Tape

Although more and more organizations are adopting disk into their backup environments, the need to maintain backup data on physical tape media will continue for the foreseeable future. Most organizations use virtual tape libraries as a way to reduce the amount of backup data going to tape, not necessarily to eliminate tape altogether.

ProtecTIER, in a Virtual Tape Library Solution, reduces offsite vaulting costs by approximately 30 percent by lengthening the online retention schedule and changing the archive to tape requirements (archive monthly, not weekly, for example).

Since ProtecTIER transparently emulates an automated tape library, all TSM tape copy, migration, and vaulting options are available. The ProtecTIER VTL and an actual tape library are attached to the TSM server either directly or through a storage area network. ProtecTIER serves as the onsite recovery repository and the tape library provides copies for offsite storage. This reduces the overall cost of the backup infrastructure and lowers operational costs and tape requirements.



There are different ways to use the ProtecTIER VTL and a real physical tape library with Tivoli Storage Manager. IBM recommends that you connect the tape library to the TSM server in parallel with VTL. The TSM server then sees two tape libraries. It thinks they are both actual tape libraries, but only one is a physical tape library, the other is a virtual tape library. The physical library is defined as the next storage pool to the ProtecTIER VTL. This allows TSM to manage the migration of data from the VTL to physical tape.

After migration, the restore path is different. In case of a restore, data must be restored from physical tape or migrate the data needed back to the VTL with the MOVE NODEDATA command.

Most customers migrate data from Virtual Tape Library Solutions to physical tape after most of the backups are done. This enables you migrate the data when the workload on the TSM server is low. Before you migrate data from the Virtual Tape Library Solution to physical tape, consider creating the copy pools, if desired, on the virtual tape library.

The combination of TSM and ProtecTIER also leads to greater flexibility and better utilization of tape media. VTL-to-tape migration can occur as a secondary process outside of the original backup window. Since the migration window is much larger than the original backup window, scheduling is simplified.

## Enhancing Progressive Incremental Backups


TSM administrators use progressive incremental backups as a way to reduce the amount of data that has to be sent and stored. Progressive incremental backups only back up new or changed data, eliminating the need to do full backups and reducing the overhead associated with other vendor's incremental or differential backups.

Progressive incremental backups require less network bandwidth and result in faster backup times and less storage required. However, they can elongate recovery times and decrease the efficiency of tape cartridge utilization. IBM uses two methodologies, called collocation and reclamation, to alleviate these issues.

## Improving Collocation

One of the only problems with progressive incremental backups is that, over time, backup data may be scattered across many tape volumes. This significantly slows restore times because it often requires many tapes to be mounted, scanned, and read. TSM uses collocation to collect all backup images of a client and store them on a minimum number of tape cartridges. This improves restore times by reducing the tapes needed. Collocation also improves restore times by enabling more restores to run concurrently. This is because the data required to restore a particular client is on its own media. In general, restoring data from a progressive incremental backup is very media-intensive, and its performance and reliability depends on the integrity of each tape required to fulfill the restore request. Collocation is designed to improve restores by minimizing the number of tape mounts associated with such restores.

A collocation operation often requires numerous tapes to be mounted and the use of several tape drives to store associated data in a common location. ProtecTIER, in a Virtual Tape Library Solution, dramatically improves collocation performance and efficiency because having the source images on disk eliminates the delays associated with loading, mounting, and scanning for data on tape. In addition, if two or more source images have been written to the same tape, the collocation process will require multiple passes over the source tape media, greatly slowing down the entire process. Since ProtecTIER technology does not suffer from these delays, the time required to complete collocation operations is dramatically improved.



The extreme flexibility of Virtual Tape Library Solutions also means the number of virtual drives and the size of the virtual tape cartridges can be defined to provide optimal backup and restore performance for the amount of data being stored and the number of clients being backed up. So even faster restore speeds can be achieved since ProtecTIER ensures that the data is written sequentially to a single piece of virtual media.

## Improving Reclamation

TSM uses reclamation to combine multiple partially used tape cartridges onto new media to create almost fully utilized tape cartridges. Reclamation increases media utilization and minimizes the number of tape cartridges needed to be transported offsite for disaster recovery purposes.

There are some problems associated with the reclamation process. In some larger environments, it may be difficult to complete the reclamation process in a given time frame. Depending on the size of the disk pools, tape volume sizes, and the number of tape drives available, the reclamation process may not complete before the next backup window begins.

Reclaiming data on virtual tape cartridges, on the other hand, is a very efficient and quick process. ProtecTIER, in Virtual Tape Library Solutions, eliminates the delays associated with loading, mounting, and positioning tapes and minimizing the time required to locate and reclaim the data. With one of these Solutions you can easily create the best cartridge size, so the reclamation threshold is optimally utilized. In addition, you can create the right amount of virtual tape drives to increase reclamation performance and minimize resource conflicts.

Another problem with the reclamation process is that it consumes a huge amount of tapes and ties up lots of system and tape resources. Many companies do not address these problems, they just continue to increase their investments in media, tape drives, and tape libraries to meet the needs of the reclamation process. A Virtual Tape Library Solution with ProtecTIER technology eliminates these problems and prevents wasteful expenditure on unnecessary tape resources.

TSM disk pools also suffer from fragmentation, which can lead to poor performance. Unfortunately, you can't use reclamation on TSM disk pools. To avoid this, some TSM administrators define file pools on disk. However, this requires the use of a file system layer that must be managed and has a negative impact on performance.


Virtual Tape Library Solutions with ProtecTIER avoid all these problems and complications. That's why most organizations will choose these Solutions to reduce the size of their disk pools, sometimes by as much as 95 percent, saving time and money, while greatly improving the overall performance of backup and recovery operations.

## Conclusion

Combining the advanced capabilities and features of TSM with the powerful performance-enhancing and cost-reducing capabilities of ProtecTIER, Virtual Tape Library Solutions by Hitachi Data Systems provide IT organizations with a powerful, cost-effective way to improve the performance, reliability, and scalability of their data protection environment.

IBM and Hitachi Data Systems have worked together to create a unique integrated data protection solution that:

- Accelerates backups to meet ever-shrinking backup windows
- Enhances ability to rapidly restore mission-critical data
- Increases reliability of backup operations
- Enables proactive verification of backup data integrity

- 
- Increases onsite recovery capacity to retain months of backup images using a small amount of storage
  - Improves backup and recovery operations without changing existing backup processes, procedures, and/or policies
  - Extends disaster recovery capabilities
  - Significantly lowers the TCO while improving data protection capabilities

With this powerful combined solution, IT organizations no longer need to choose between optimizing for backups and optimizing for restores. IBM and Hitachi Data Systems make the entire backup and recovery environment faster, simpler, more reliable, verifiable, and more cost-effective, enabling administrators to meet increasingly stringent service levels and reduce overall costs.



---

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

Contact Information: 1 408 970 1000 [www.hds.com](http://www.hds.com) / [info@hds.com](mailto: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](mailto:info@hds.com)

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

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

TrueCopy is a registered trademark of Hitachi Data Systems Corporation.

IBM and Tivoli are registered trademarks of International Business Machines Corporation.

ProtecTIER, HyperFactor, and Diligent are trademarks of Diligent Technologies 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](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-253-00 LKW May 2007