




Virtual Tape Library

Solutions Brief

Improving the Cost, Reliability, and Responsiveness of Your Backup and Recovery Infrastructure

Partner Beyond Technology





Executive Summary

Demand for greater application availability and unprecedented growth in storage capacity are two of the biggest challenges facing enterprise IT today. These two trends are affecting many areas of the technology environment, but few are showing the stress exhibited by the data protection infrastructure.

Magnetic tape, the predominant storage medium for backing up application data, is now seen as a major stumbling block to desperately needed improvements in backup and recovery processing. As the amount of data to be protected escalates and business user demand for application availability increases, the physical, mechanical nature of tape is preventing innovative new approaches to backup processing.

Virtual Tape Libraries (VTLs) offer a new way of thinking about enterprise data protection. Providing a virtual interface that mimics existing tape resources while writing backup data to disk, the VTL addresses the limitations of tape-based processing without changing the basic paradigm of enterprise data protection.

The VTL reduces backup time, speeds recovery from failures, improves reliability, reduces staffing requirements, enables automated remote operations, and provides for centralized backup, recovery, and business continuity processing. And, using innovative data de-duplication technology and seamlessly integrating into an existing infrastructure, the VTL is able to dramatically lower the total cost of ownership of the data protection infrastructure, as shown in Figure 1.

The Problem with Tape-based Backups

Organizations have relied on magnetic tape as a backup storage medium since the earliest days of IT. Inexpensive and portable, tape storage has been the workhorse of IT operations, providing a cost-effective medium for backup data and allowing copies of critical information to be carried offsite to be stored at a safe location in case of emergency.

Despite constant improvement by technology vendors, tape has many well-known drawbacks.

Failure

Magnetic tape is a mechanical technology, with cartridges, tape drives, libraries, and silos having many moving parts, all of which are subject to failure through wear and tear. The mechanical nature of tape means that it is inherently unreliable. Backup jobs that write to tape experience a 20-to-30 percent first-time-run failure rate, with failures almost exclusively linked to mechanical errors.

Backup failure has consequences for the data preparedness of an organization. When a backup job fails, operators must choose between rerunning the job, potentially extending the backup window and restricting application availability, or canceling the scheduled run, leaving application data unprotected until the backup run the following evening.

The Backup Window

To guarantee the integrity of a backup, all application writes must be halted while data is written to tape. For most users, applications are effectively unavailable while the backup job is running.

The window of time when backups are run is continually under pressure. Business users, demanding more application availability, want a shorter backup window. However, longer backup run times, due to increased data volume, are continually extending the length of the window.

The size of the backup window is a function of how much data must be backed up, how many tape drives are available to support concurrently running backup jobs, and how fast the data can be written to tape. Vendors of tape drives and backup software have made significant improvements in tape write speed using techniques like streaming. However, the sequential nature of tape access and the need to stop and reposition the tape periodically can slow backup processing.

Capital Expenditure

The number of tape drives, tape cartridges, and tape libraries and silos housing the cartridges provides a measure of the size of a data protection infrastructure. As demand for storage grows, IT planners must ensure the infrastructure has sufficient capacity to accommodate an increase in backups.

Increasing the capacity of a tape infrastructure is expensive. Although big-ticket items like silos and tape libraries provide room for growth, scalability is not incremental. If the

capacity of existing equipment is exhausted, moving to the next level involves considerable capital investment. Not only must new equipment be purchased, but additional room in the data center must also be found for the new resources.

Application Recovery Point and Recovery Time Objectives

Backup processing represents one half of the data protection equation. The recovery process enables a known good copy of application data to be substituted for the primary copy, if an unforeseen corruption occurs. During recovery, all application access is halted while the data is restored from a backup tape. This makes the speed of the recovery critical for many important applications.

Business continuity planners measure the recovery needs of an application using two key metrics: the recovery point objective (RPO), and recovery time objective (RTO). RPO indicates the point at which a recovery process must return application data to, and is usually thought of as a measure of how much application data an organization can tolerate losing before suffering financial penalties.

RTO signals how quickly an application must be available to users after an unforeseen outage. The speed with which data can be recovered from tape is a significant factor in meeting application RTO service levels.

Recovering application data from tape is cumbersome and slow. When a data corruption occurs, the appropriate backup tape must be located in the library and

mounted to a tape drive. The recovery application must then search sequentially through the tape to find the right data to restore. Each delay in recovering data corresponds to downtime for business applications.

Tape Management

Tape is not the most efficient storage medium in the data center. Although tremendous advances have been made in increasing the storage capacity of each physical cartridge, most tapes are not fully used.

Organizations generally deploy a grandfather-father-son (GFS) media rotation schedule for backup tapes. In a typical GFS configuration, incremental backups are run daily (five sets of tape cartridges), full backups are run at the weekend (52 per year), the last full backup each month is retained as a monthly (12 per year), and the last monthly per year is saved as an annual backup (retained for seven years). Using the GFS rule, over the course of a year, every terabyte (TB) of primary disk storage can require up to 25TB of tape, depending on retention periods. Given that tapes are rarely fully used, this represents an enormous amount of unused storage capacity.

Using a GFS configuration, an organization with 30TB of primary storage must plan on having 750TB of tape capacity—approximately 9000 LTO-1 cartridges. Planning for retirement and replenishment of cartridges and offsite transport and storage is a considerable logistical and economic challenge.

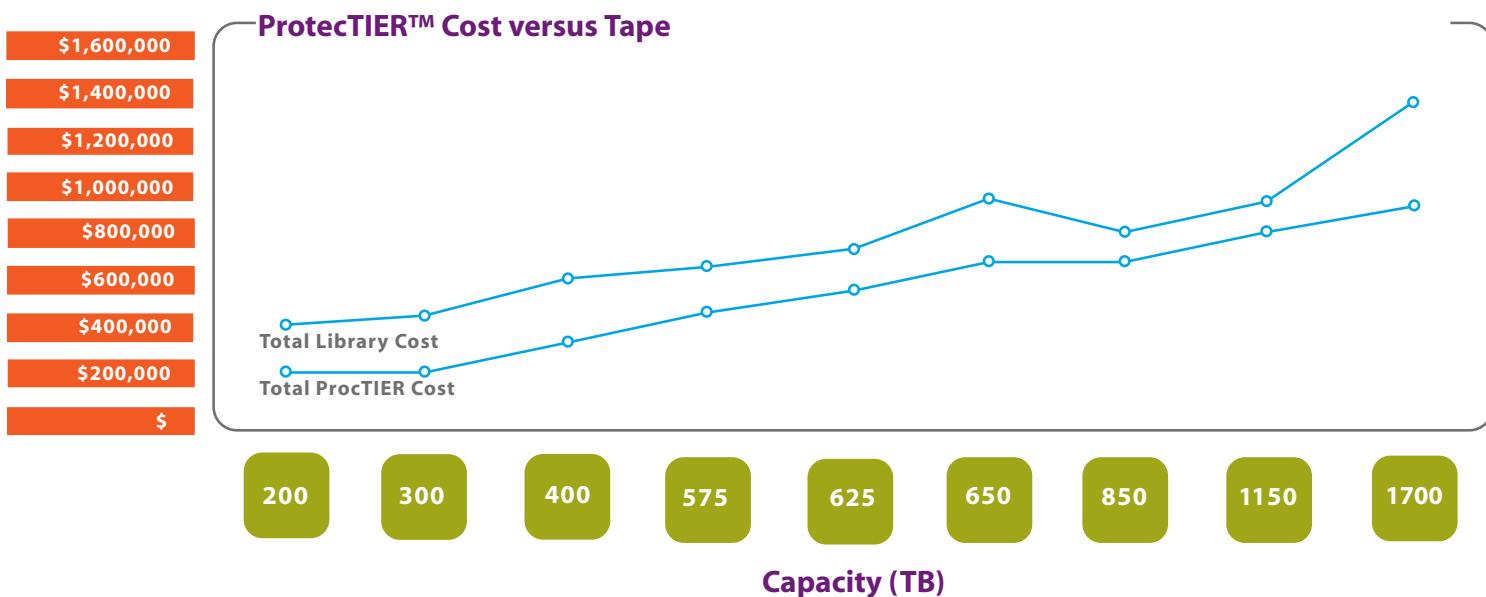


Figure 1: The cost of traditional tape versus a Virtual Tape Library Solution by Hitachi Data Systems, employing ProtecTIER technology.

Virtual Tape Library Solutions

Virtual Tape Library Solutions from Hitachi Data Systems eliminate the problems associated with conventional tape-based backup processing. Seamlessly simulating tape operations while writing backup data to disk, these solutions integrate into an existing backup and recovery infrastructure without requiring any change to processes, procedures, or policies.

Writing backup data to disk, rather than tape, improves the reliability of backup and recovery processing, satisfies demanding application RTO requirements, provides cost-effective scalability unmatched by any tape-based alternative, and significantly alters the economics of backup and recovery.

Improve Reliability

Although no system is immune to failure, vendors of disk storage have had many years experience developing solutions that safeguard critical data from loss in the event of a system outage. RAID storage configurations, using multiple disks to ensure redundancy, provide availability to satisfy the most stringent demands of online applications.

Virtual Tape Library Solutions by Hitachi Data Systems substitute Hitachi RAID storage for tape in a backup environment. This transforms the reliability of the data protection infrastructure and virtually eliminates backup job failures and costly reruns.

Reduce Backup Window

Improving reliability provides much needed first-time-run stability for backup jobs, giving IT operations greater control over the backup window. This stability enables IT executives to

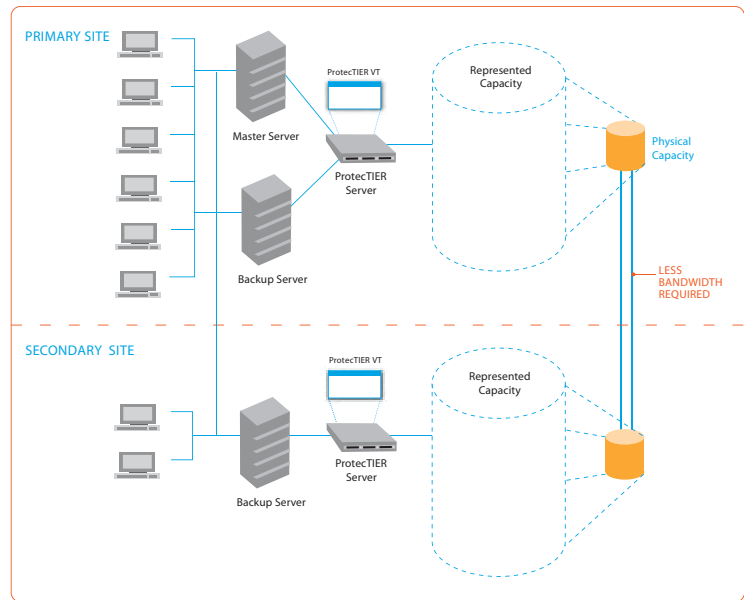


Figure 2: Diligent™ Technologies Corporation's ProtecTIER, with HyperFactor™ technology, reduces bandwidth needs for replication.

provide business leaders with hard and fast service-level guarantees for availability and for application RPOs and RTOs.

Capable of receiving data at hundreds of megabytes per second, and eliminating the stop, start, and reposition process of tape, the VTL is able to significantly increase the speed of backup processing. This provides scheduling consistency for backup operations, and faster processing to further reduce the length of the backup window.

Increase Availability

In addition to faster writes, the random access nature of disk also provides substantial improvements in recovery time. By eliminating the need to locate and mount tapes and then search sequentially for files to restore, disk-based recoveries become almost instantaneous, resulting in improvements in meeting application RTO service levels.

With a faster and more reliable backup process, IT operations can choose to run backups more frequently. This can have a significant impact on application RPO by reducing the potential for data loss after a system outage.

Automate Business Resilience

Disk-based application data is more portable than tape. Fast storage networks now routinely connect local and remote enterprise data centers, allowing replication technologies to automate the secure transmission of critical disk-based backup data for storage at a remote location. Replicating backup data between sites eliminates the previously manual and error-prone process of carrying tape cartridges offsite for transport to a business continuity facility.

Leveraging data replication technologies enables the VTL to further improve the resilience of the organization after unforeseen outages. By replicating VTL contents to a corresponding library at a remote location, as shown in Figure 2, application data is available immediately for recovery should anything disrupt access to the primary data center.

Incremental Scalability

With constantly escalating demand for storage capacity, the scalability of the data protection infrastructure is a critical concern to IT planners. VTL solutions provide significant scalability improvements compared to traditional tape.



Figure 3: The multilevel framework of Application Optimized Storage™ solutions from Hitachi Data Systems delivers the components necessary to optimize storage and meet business needs for efficient operations.

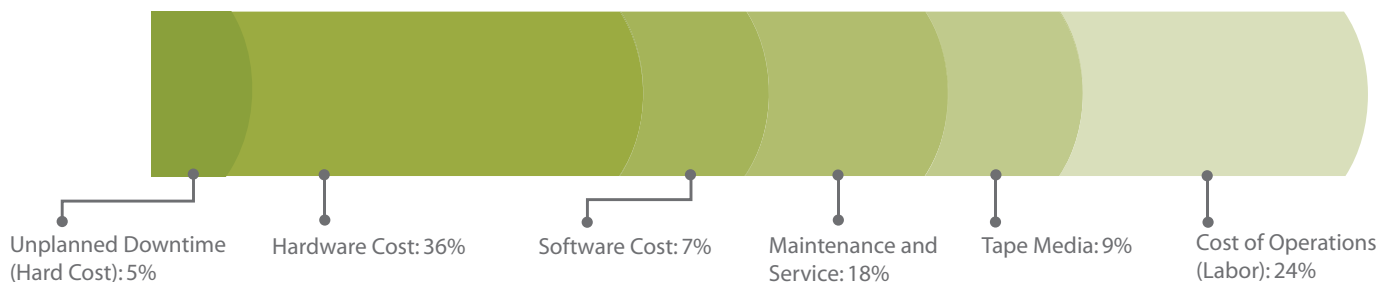


Figure 4: Projected three-year TCO for traditional tape-based backup and recovery infrastructure components.

Because the VTL's simulated tape resources are virtualized, there is no additional cost to allocate more cartridge or tape drive capacity. Each node of the VTL can support up to 200Mbit/sec of input data. Additional nodes and storage provide smooth incremental expansion to accommodate increased demand for backups.

Redefining the Economics of Data Protection

Although vendors of tape storage systems have made dramatic improvements in reliability, speed, and capacity, the primary advantage of tape has always been economic. However, innovations in disk storage have steadily closed the gap and disk now represents a cost-effective backup storage medium with none of the problems traditionally associated with tape.

Furthering the reduction in cost, technologies that comb through the input backup stream and eliminate duplicate data, already stored on the VTL's disk, can significantly lower storage requirements. HyperFactor™ de-duplication technology from Diligent Technologies can radically alter the backup and recovery infrastructure total cost of ownership (TCO) equation, making disk-to-disk backups substantially more cost-effective than tape.

Application Optimized Storage™ Solutions: Aligning Business and IT Objectives

The Hitachi Data Systems Application Optimized Storage™ solutions optimize the cost of the storage infrastructure while guaranteeing applications access to the functionality they need. An Application

Optimized Storage solution is founded on a tiered storage infrastructure, in which the capacity used by enterprise data is hosted on storage that meets the exact performance, availability, functionality, and cost requirements of the application's users, as shown in Figure 3.

Virtualization is a key enabling technology of Application Optimized Storage solutions. Providing a common interface to capacity from different vendors, virtualization allows IT managers to provision resources to satisfy the diverse needs of enterprise applications without introducing complexity into the infrastructure.

VTL solutions are an essential component of an Application Optimized Storage infrastructure. By increasing the reliability and speed of the backup process, reducing the impact of backups on application availability, and delivering faster recovery from failures, the VTL provides a critical tier in a multitiered storage infrastructure. Business continuity is critical to successful business operations and Virtual Tape Library Solutions by Hitachi Data Systems improve an organization's ability to protect critical application data from loss.

ProtectTIER: Reducing the Total Cost of Ownership

Faced with a projected 50 percent annual growth in storage capacity, a need to cut the ballooning cost of data protection, and business user demand to reduce downtime, the IT group at a large, well-known retailer began to evaluate disk as an alternative backup target to tape. Over three years, the retailer's primary application storage capacity

was expected to grow from 20TB to 67.5TB, and this was putting unprecedented strain on the backup and recovery infrastructure.

Averaging five large recoveries and 15 smaller recoveries per month, the two fulltime IT staff dedicated to backup and recovery were hard pressed to maintain the existing environment, which consisted of:

- ❖ 12 LTO-1 tape drives in a library
- ❖ GFS backup rotation schedule with tapes sent offsite at the end of each week
- ❖ 6200 tape cartridges

When taking into account all incremental costs—hardware, software, maintenance, service, media consumption, administrative labor, and unplanned downtime—the annual TCO of the backup and recovery environment was projected to be US\$5.5 million over three years. (See Figure 4.)

The ProtecTIER Advantage

Using a Virtual Tape Library Solution by Hitachi Data Systems, which incorporates ProtecTIER, the retailer was able to show that projected three-year TCO for the backup and recovery environment would drop to US\$2.4 million: a savings of over US\$3 million. (See Figure 5.)

The ProtecTIER savings come from a wide range of influences on the backup and recovery infrastructure. Hardware, labor, and maintenance and service costs, the top three categories for a tape-based infrastructure, are substantially reduced. And, because ProtecTIER uses innovative HyperFactor de-duplication technology, backup

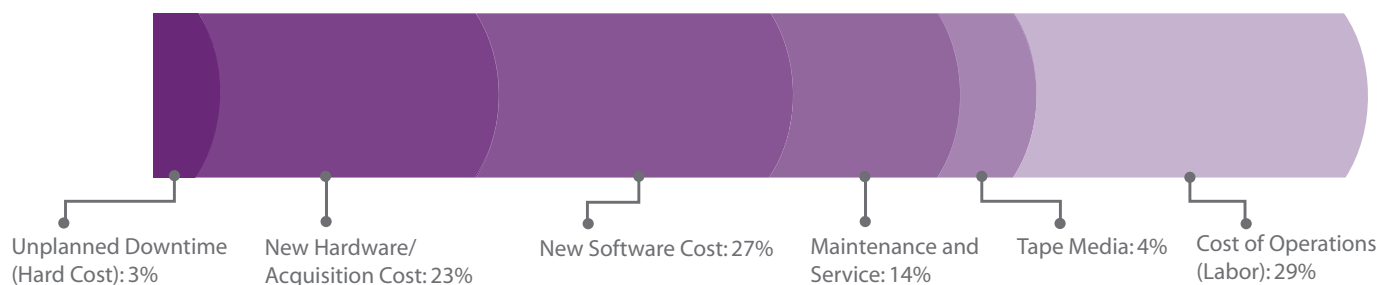


Figure 5: Projected three-year TCO for ProtecTIER backup and recovery infrastructure components.

and recovery storage requirements are dramatically reduced. After four weeks, ProtecTIER is able to reduce 30TB of conventional tape cartridge backup capacity to 19TB of VTL storage. After 17 weeks, as the HyperFactor algorithms recognize more duplicate data, ProtecTIER requires only 24TB of VTL storage to host the equivalent of 287TB of tape backup cartridges. The Virtual Tape Library Solutions from

Hitachi Data systems, using Diligent Technologies ProtecTIER software, eliminate the need to purchase additional tape library expansion units, tape drives, and tape media. And, because these VTLs are hosted on Hitachi storage systems, the retailer can make use of the innovative replication technologies from Hitachi Data Systems to further reduce tape media and services costs, and enhance the business resilience of the organization.

Deploying a Virtual Tape Library Solution

This solutions brief provides a first step on the road to understanding the implications and challenges of integrating a virtual tape library solution into your data protection infrastructure. The information in this brief illustrates the substantial benefits that accrue from switching to a virtual tape library solution. This brief will help you to build a business case justifying implementation of a virtual tape library in your environment.

Although reading is a necessary first step in understanding the ramifications of virtual tape, Hitachi Data Systems strongly recommends that customers engage our Global Solution Services group before they seriously undertake a review of their long-term data protection strategy. Global Solution Services consultants have substantial experience implementing virtual tape library solutions and can help you to assess, design, plan, and implement the optimal solution to meet your business and application requirements. The Hitachi Data Systems Global Solution Services consultants can create, deliver, and integrate a custom solution that enables you to meet RPO and RTO objectives for mission-critical data and provide a balance between increasing data capacity and shrinking IT budgets.

To learn more about how Hitachi Data Systems can help with your virtual tape library plans and to read more about Virtual Tape Library Solutions by Hitachi Data Systems, please visit www.hds.com or call Hitachi Data Systems at 888 234 5601, ext. 950, to explore an engagement that will result in the optimal solution for your data protection needs. (Hitachi Data Systems Channel partners should contact their Channel managers for information.)

Virtual Tape Library Solutions by Hitachi Data Systems

Virtual Tape Library Solutions by Hitachi Data Systems are a highly scalable set of enterprise-class solutions featuring software applications from Diligent Technologies. These solutions seamlessly integrate into an existing infrastructure, without requiring changes to data protection practices, policies, or procedures, and allow backup applications to use fast and reliable disk as a backup medium.

The Virtual Tape Library Solutions use a combination of software, server, and Fibre Channel-attached storage systems to emulate one or more tape libraries, as shown in Figure 6. Each solution writes backup data to disk, giving users dramatic improvements in speed, reliability, recoverability, and cost of backup and recovery operations.

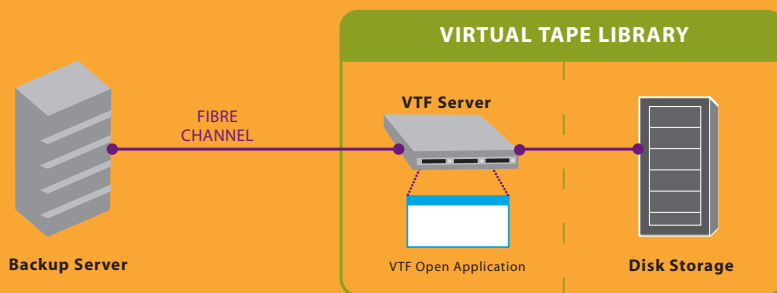


Figure 6: VTF™ Open basic configuration.

VTF™ Open

VTF™ Open is an enterprise-class tape emulation software solution for open systems running on a dedicated Linux server, and supporting clustering for effortless scalability. Integrating seamlessly with existing backup and recovery software, including Symantec NetBackup, Tivoli Storage Manager, Hitachi Data Protection Suite, and CommVault Galaxy, VTF Open supports up to 50 servers and more than 5TB of disk-based primary data and an unlimited number of virtual tape cartridges.

VTF™ Mainframe

VTF™ Mainframe is a scalable software solution that emulates tape libraries for mainframe-based backup and recovery applications. Fully integrated with any IBM® z/OS® environment, VTF Mainframe does not require any changes to existing JCL or backup and recovery applications, and supports more than 5TB of disk-based primary data.

ProtecTIER™ VT

ProtecTIER™ VT is the industry's first enterprise-class data protection platform that reduces backup storage capacity requirements by up to 25 times or more, without disrupting existing infrastructures. Powered by a revolutionary new algorithm called HyperFactor™, the ProtecTIER platform eliminates redundancy by only backing up data that has not been previously stored, enhancing the physical storage capacity of the disk storage. ProtecTIER software runs on a standard enterprise Linux server, with clustering support for scalability, and connects to a Fibre Channel-attached storage system.

Radically Reduce Infrastructure Costs

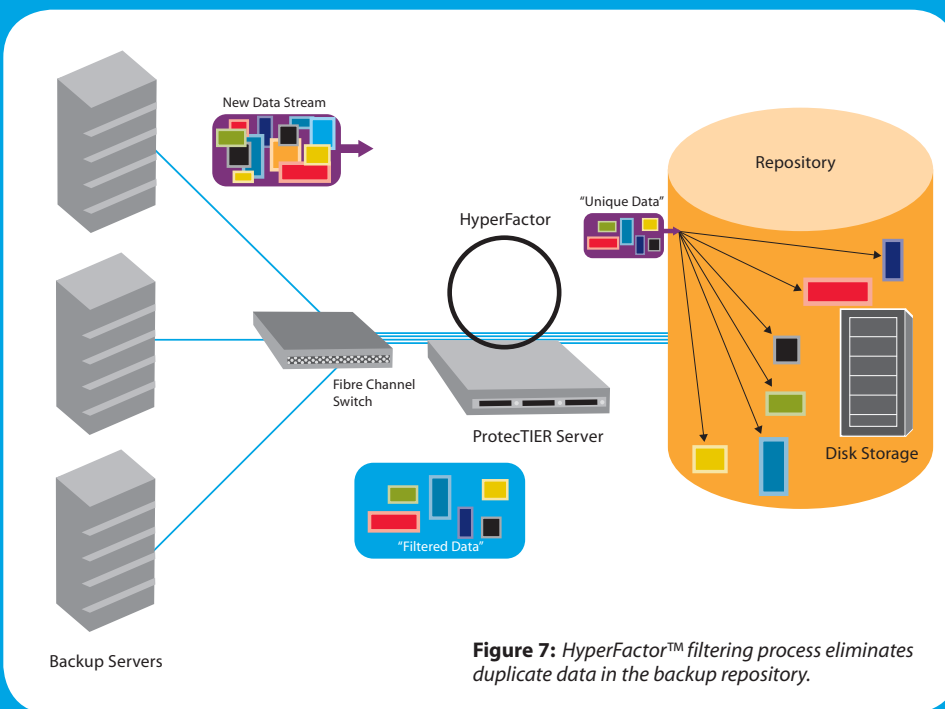


Figure 7: HyperFactor™ filtering process eliminates duplicate data in the backup repository.

The ProtecTIER VTL solution is powered by the Diligent Technologies HyperFactor algorithm: a breakthrough in data factoring technology, as shown in Figure 7. HyperFactor reduces backup storage needs by up to 25 times or more, compared to conventional tape-based backup methods. And, most critically, HyperFactor achieves these dramatic savings without compromising backup data integrity.

Using a series of algorithms to efficiently map backup data contents, HyperFactor eliminates duplicate information from the backup data stream as it is written to disk. A single instance store provides a repository for common elements. Only

new data and pointers to the duplicate information are written to the storage system.

HyperFactor uses a highly efficient RAM-based index capable of mapping 1PB of physical storage into 4GB of memory—representing a 25,000-to-1 storage-to-memory ratio. The RAM index eliminates the need for disk access and costly I/O during factoring, allowing ProtecTIER to maintain high backup data throughput rates.

HyperFactor is up to 1,000 times more efficient than competing factoring techniques that use hash-based data coalescence. This results in better performance and savings in memory on the ProtecTIER server. Over time, the combination of ProtecTIER and HyperFactor reduce physical backup storage capacity by up to 25 times or more, dramatically altering the economics of protecting enterprise data.

Complementary Solutions

These solutions complement the Virtual Tape Library Solutions by Hitachi Data Systems.

Hitachi ShadowImage™ In-System Replication Software

Hitachi ShadowImage™ In-System Replication software is business continuance software that provides a host-independent data replication solution for duplicating any customer-accessible data within a single Hitachi storage system.

Hitachi TrueCopy™ Remote Replication Software

Hitachi TrueCopy™ Remote Replication software is business continuance software that provides a continuous, host-independent data replication solution for duplicating local data onto one or more physically separated Hitachi storage systems.

Hitachi Universal Replicator Software

The Hitachi Universal Replicator software delivers asynchronous data replication for the Hitachi TagmaStore® Universal Storage Platform and Network Storage Controller. It introduces innovative technology to support advanced business continuity requirements.

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

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.

TagmaStore is a registered trademark and Application Optimized Storage, ShadowImage, and TrueCopy are trademarks of Hitachi Data Systems Corporation.

Diligent, VTF, ProtecTIER, and HyperFactor are trademarks of Diligent Technologies 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 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.
DISK-603-00 July 2006

