

Replication Solutions for Hitachi Modular Storage

White Paper

By Thomas Neidhardt

December 2008



Executive Summary

Protecting data and information critical for the operation of a business, whether it is a large enterprise or a small to medium sized business, is becoming increasingly important. Large enterprises have a staff of IT professionals that support their storage environments and are typically responsible for ensuring that the data will be available in the case of an operational error or a catastrophic event. For smaller businesses, managing the storage environment and ensuring that critical data will always be available is more difficult as they often have a minimal or non-existent IT staff. As a result, data is often not regularly backed up or replicated to another system or site for protection. The perceived cost, complexity and/or time involved are often overwhelming, resulting in delays in addressing the need or just ignoring the potential problem.

Any comprehensive storage solution must include methodologies and support for protecting data while facilitating continuous business operations in the face of localized operational problems or large-scale catastrophic events. These solutions range from providing local copies of data to completely redundant environments at recovery sites located a great distance from the day-to-day operational environment.

A well thought out and implemented data protection plan can help ensure the ability of business continuance in the case of local operational problems or large-scale catastrophic events. Such a plan can reduce the time to recover and restart business operations. It can also improve operational efficiency and productivity by providing access to accurate, current copies of production data for use in data mining operations, application development testing and other non-production operations.

Hitachi Data Systems provides comprehensive storage solutions for any storage need, large or small. In this paper we will discuss modular storage solutions from Hitachi Data Systems for the small to medium business.

For best results use Acrobat Reader 8.0.



Contents

Introduction	1
Local Replication	1
Hitachi Copy-on-Write Snapshot Software	2
Hitachi ShadowImage® Replication Software	3
Remote Replication.....	4
Hitachi TrueCopy® Synchronous Software.....	4
Hitachi TrueCopy Extended Distance Software	5
Hitachi Protection Manager Software	8
Protection Manager Software with Microsoft Exchange	8
Protection Manager Software with Microsoft SQL Server	9
Benefits	9
Replication Management.....	9
Summary	10

Replication Solutions for Hitachi Modular Storage

White Paper

By Thomas Neidhardt

Introduction

The first step in developing a plan to protect a company's data is to understand the unique requirements for that business. Not all data is critical to the continued operation of a business; however, there is usually critical information, such as transaction data, customer files, inventory data, financial records and business or manufacturing process data, whose availability can be the difference between continued operations of a business or its failure.

It may be impossible to recreate some data while other data may be easy to recreate. Copies of data may need to be readily available in case there is a local equipment failure or operational error that renders the active data unusable. These copies should be located within the data center or primary business location. Additional copies may be required to be kept offsite and only used in case of a major disruptive event where there is time to retrieve the data from a remote location before continuing operation. The third situation may be that data must be immediately available at a remote site. Current regulations require that some industries must be operational after a very short time even in the case of a catastrophic event. In this case, the copies of the data must be current and available at a remote site that is outside of the catastrophic area.

As a starting point, it is necessary to have a highly available storage environment that provides nondisruptive data replication services for both local replication as well as replication to a remote storage system. Data replication should be transparent to the host systems, including the application and operating system.

Hitachi Data Systems understands these varied requirements and provides the products, services and solutions to address them. The Hitachi Data Systems Global Solution Services (GSS) team includes experienced business continuity professionals who work to identify an organization's unique requirements, including budget and staffing issues, and develop a plan to address them.

Local Replication

Hitachi Data Systems provides modular storage products that address the need for local replication of data on the modular storage family of products, which include:

- Hitachi Adaptable Modular Storage 2500
- Hitachi Adaptable Modular Storage 2300
- Hitachi Adaptable Modular Storage 2100
- Hitachi Adaptable Modular Storage 1000
- Hitachi Adaptable Modular Storage 500
- Hitachi Adaptable Modular Storage 200
- Hitachi Workgroup Modular Storage 100

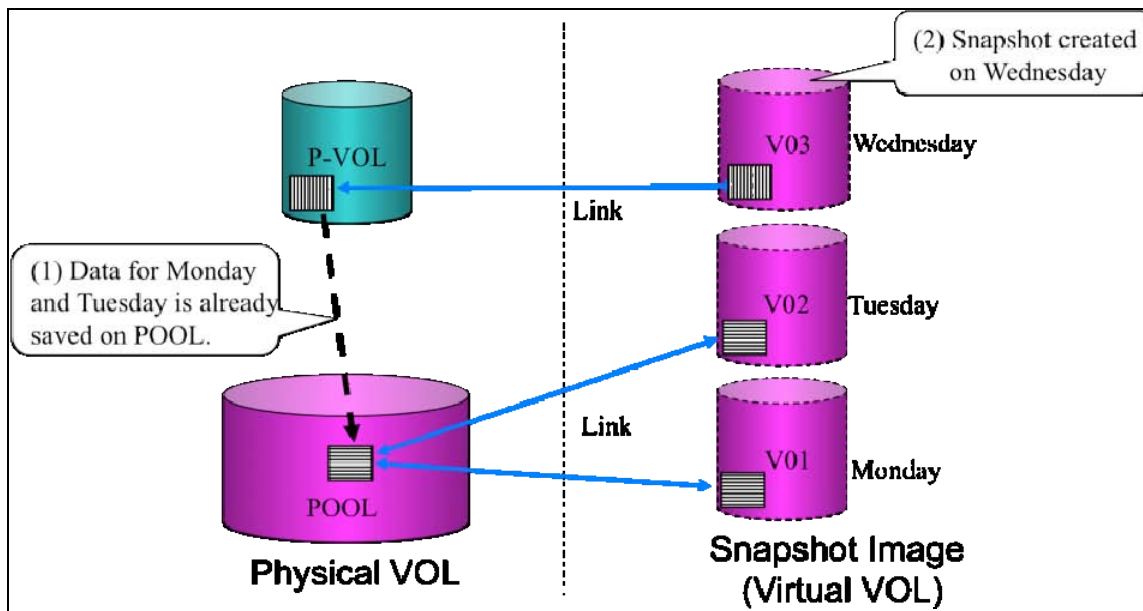
Hitachi Copy-on-Write Snapshot Software

Hitachi Copy-on-Write Snapshot software is an optional storage feature, which can nondisruptively create up to 32 point-in-time snapshots of any data volume within the Hitachi storage system. Snapshot technology creates images of the data by only copying changed data blocks to a storage pool (see Figure 1). Saving just the changed data significantly reduces the amount of physical storage required for each copy. These snapshot volumes are fully compatible for reading and writing by other hosts and can be used for rapid data restores, application testing and development, data mining and warehousing, and nondisruptive backup and maintenance procedures.

When a snapshot is created, the storage system creates a virtual volume (V-Vol) that, to the host, looks identical to the primary volume (P-Vol). In reality, it is just a table that maps all the P-Vol data blocks. Before any writes to the P-Vol take place, all of the entries for the data blocks point to the P-Vol. When a write to the P-Vol occurs, the data block that is being changed is copied to the storage pool and the V-Vol table entry for that data block is updated to point to the block in the storage pool. When a read is issued against the V-Vol, the storage system uses the V-Vol table entries to either read the data from the primary volume if the data block has not been changed or read the data block from the storage pool if the block has been changed.


If a new snapshot image (V-Vol) is created for the same volume, a new table is created and is managed the same way as described above. Each V-Vol snapshot is time stamped, so if it is necessary to recover a volume, the snapshot taken just before the recover point can be used to restore the P-Vol to the point in time of that snapshot.

Figure 1. Hitachi Copy-on-Write Snapshot Example



In Figure 1, a snapshot was taken on Monday and another one was taken on Tuesday, with the corresponding changes saved in the storage pool. When the snapshot for Wednesday is created, a new table is created and initially points to the P-Vol for all the data blocks on the volume.

Hitachi modular storage facilitates great flexibility for recovering data to a specific point in time. Using the Adaptable Modular Storage 200, 500 or 1000 and the Workgroup Modular Storage 100, up to 15 generations of snapshots — of hundreds of volumes — can be saved. The Hitachi Adaptable Modular Storage 2000 family extends the maximum number of snapshots per volume to 32.



The benefits of Copy-on-Write Snapshot technology include:

- Having an instantaneous snapshot of the primary volume with no initial copy required
- Minimizing the physical disk space required since only the data that has been changed is saved to the pool
- Reducing the backup window because once the snapshot is created, it can be used to make the point-in-time backup and production does not have to be stopped during the backup
- Quickly restoring lost or corrupted data after an operational or logical error; with multiple snapshots, restore from the most current non-corrupted snapshot is easy
- Taking multiple synchronous snapshots of multiple volumes (a consistency group) so that applications using multiple volumes can have a consistent snapshot of the application

Hitachi ShadowImage® Replication Software

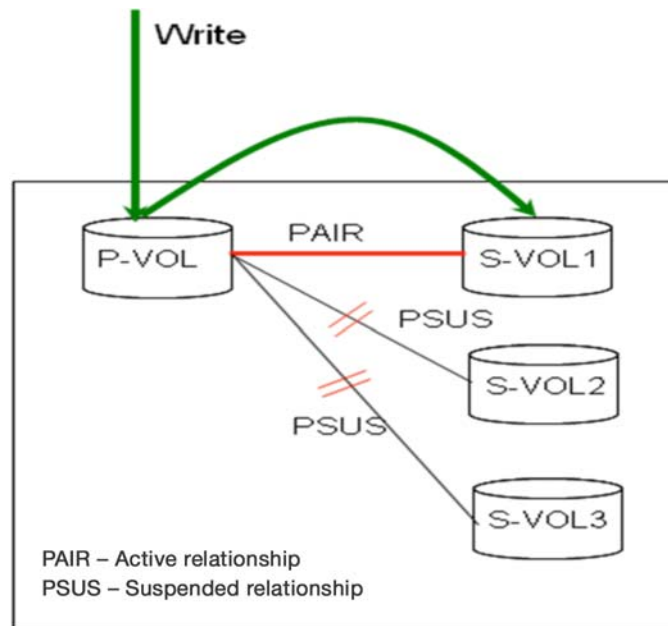
Hitachi ShadowImage® Replication software is a key storage feature within the Hitachi Data Systems Business Continuity portfolio. It provides the ability to create clones (point-in-time copies) of data with minimal impact to production (a clone is an identical full copy of the primary volume). ShadowImage can create multiple clones of a P-Vol, each of which can be used for a different purpose.

ShadowImage creates RAID-protected clones of the data within the same storage system. Like Copy-on-Write Snapshot software, ShadowImage can create consistent groups of interrelated application volumes, such as critical databases, and their associated logs are also kept consistent. This can be critical for quick recovery of an application after a disaster or operational error as well as application testing.

When a ShadowImage volume is created and the relationship between the primary volume (P-Vol) and the clone or secondary volume (S-Vol) is created, the data from the primary volume is copied to the S-Vol; this is called the initial copy. Once the initial copy completes, the two volumes are said to be in a PAIR relationship; each write to the P-Vol is also copied to the S-Vol (see Figure 2). While the relationship is in a PAIR status, the data on the S-Vol is not available for use by other applications or hosts. However, the relationship can be split — in other words, suspended or broken — and the S-Vol can then be used by another application or host. If the relationship is suspended, not broken, the storage system keeps track of any changes that have been made to tracks on the primary volume. When the PAIR relationship is reactivated, any changes made to the primary volume are copied to the secondary volume; and when all the changed tracks have been copied, the volumes are resynchronized. If the relationship is broken — in other words, split but not suspended — changes to the volumes are not tracked. When the relationship is re-established, an initial copy must again take place.

For the Hitachi Adaptable Modular Storage 2000 family, there can be up to eight S-Vols per P-Vol, but only one can be active at a time. For other Hitachi modular systems, the maximum number of S-Vols per P-Vol is three.

Figure 2. Hitachi ShadowImage Replication Software Example



The ability to test or debug applications with data that is a current replica of what is being used in production is critical. ShadowImage clones can also be used as the source for offline backup to tape. This can eliminate the issues associated with a backup window that is never long enough to complete the required backups.

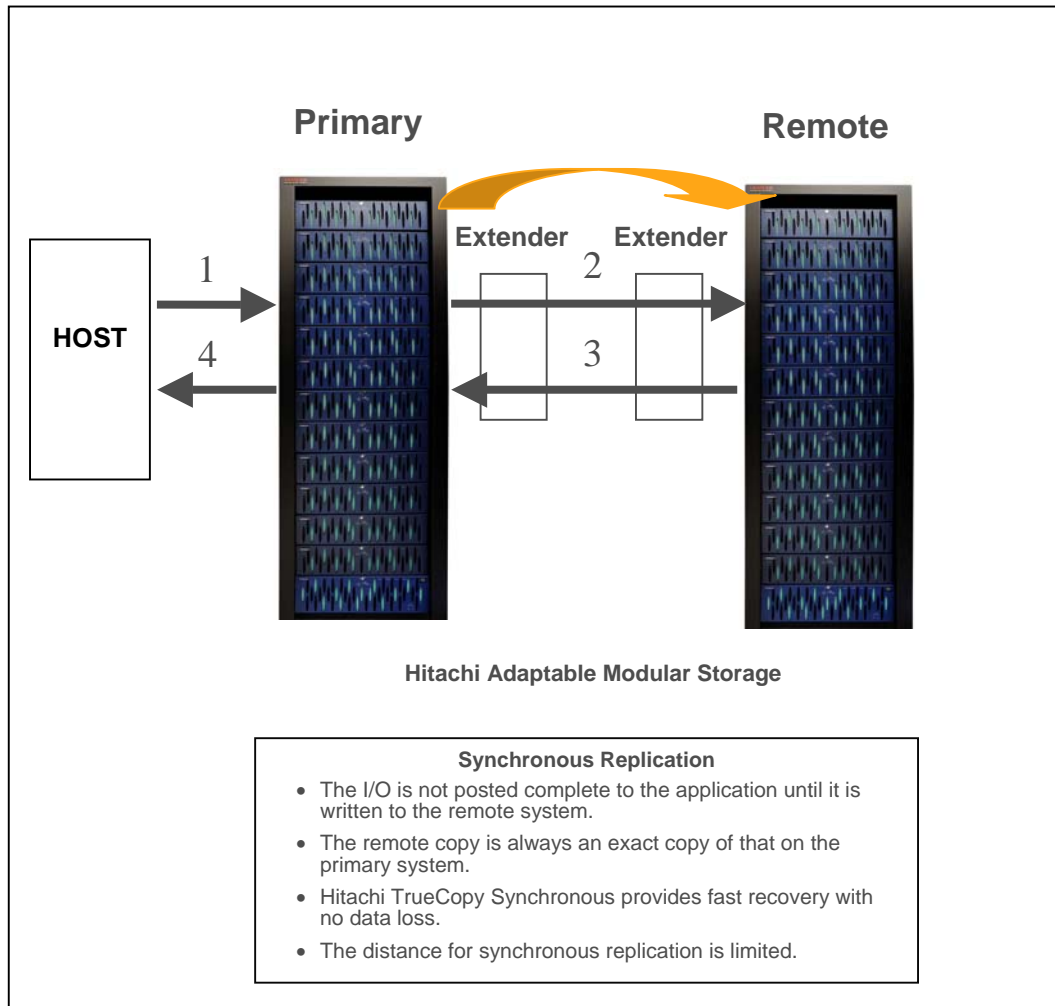
ShadowImage software provides the ability to create multiple copies/clones of data with no impact to production volumes, and without the use of host processing cycles. Once the initial copy is complete, these clones can immediately be made available for use by other applications. These clones can be used as a golden copy for logical as well as physical disaster recovery. This means that if a data volume is corrupted due to logical or physical error, administrators can use the golden copy to quickly restore back to a point in time via resynchronization of the P-Vol from a clone or by mounting a clone as the P-Vol. Administrators can also use a clone for testing or debugging of applications. If there are multiple clones, different application testing groups can each have their own copy to work with. If an application fails, debugging can proceed using an exact copy of the data that was in use when the error occurred, frequently accelerating the debugging process.

Remote Replication

Hitachi TrueCopy[®] Synchronous Software

Synchronous replication of data to a remote system provides an exact copy of the data at the primary storage system on the remote one (see Figure 3). Hitachi TrueCopy[®] Synchronous software for modular storage provides a host- and network-independent data replication solution over metropolitan distances. Due to the synchronous nature of this solution, each write to primary storage is also written to a remote direct access storage device (DASD) by TrueCopy Synchronous software before the write complete acknowledgement is sent back to the host. This means that both storage systems have the same copy of the data and are completely in sync. TrueCopy Synchronous can be used to replicate data across distances up to 300 kilometers with no problem or significant impact to the host application.

Figure 3. Hitachi TrueCopy Synchronous Software



It is important that adequate bandwidth be provisioned to move each data record with no delay in transmission. Delay introduced through insufficient bandwidth can impact the host application by delaying the write complete acknowledgement to the application. With synchronous replication, the link between the two sites is critical. If the link is lost, the administrator must decide whether to suspend processing until the link is re-established or to continue processing without the data being replicated to the remote site and risk not having an exact copy of the data at the remote site to be used in case of a major outage at the primary site.

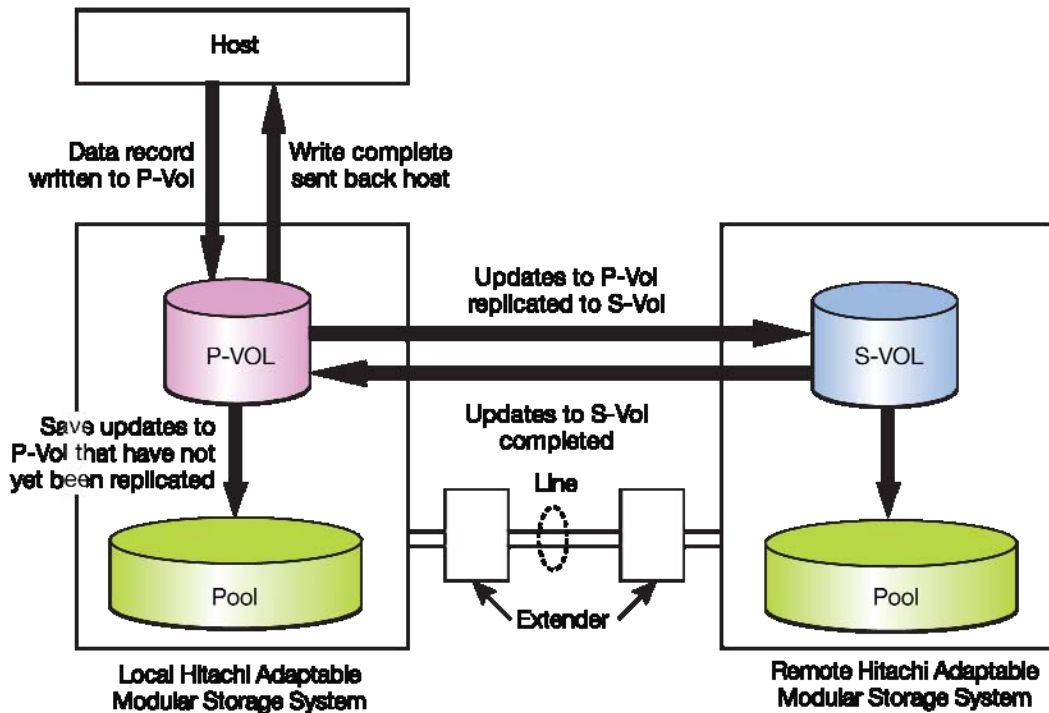
Hitachi TrueCopy Extended Distance Software

Data replication to a remote storage system requires planning and decisions on how far away the remote system should be. While synchronous replication guarantees that the data at the remote site is an exact copy of the data at the primary location, there is a limit to how far from the primary system the remote system can be located. The latency introduced by moving the data over distance increases directly with the distance between sites. The farther the remote site is from the primary site, the longer it takes to transfer the data and get an acknowledgement that the data has been successfully received at the remote site. This delay or latency can impact the processing of the primary system because application processing at the host is suspended until the acknowledgement from the remote system is received. It can extend the response time to users and can even cause the host application to time out if the delay is too great.

Asynchronous replication eliminates the problem of delaying the application since the write complete from the primary system is sent back to the host immediately while the write to the remote system is handled separately. While this eliminates the problem of the impact on the host, it does introduce a delay between when the data is written at the primary system and when it is written at the remote system. This delay means that if a problem occurs at the primary site, data records could be waiting to be transmitted or could be enroute to the remote site (being transmitted) and not yet written to the remote disk. An outage or disaster means these records would be lost and the data at the remote system would be slightly out of sync with the primary one. This difference caused by latency is typically quite small and can be managed, and the problem of any impact on the host application regardless of distance is mitigated.


Hitachi TrueCopy Extended Distance software, an asynchronous replication solution, provides a bi-directional remote data protection solution that copies data over any distance without interrupting the application, and provides basic failover and recovery capabilities for Hitachi Adaptable Modular Storage systems (2500, 2300, 2100, 1000 and 500). Because it is bi-directional, the target storage system can also be a source for replication, allowing the organization to protect critical data at both locations by having it replicated to the other location (see Figure 4). It is a flexible point-in-time copy solution that allows administrators to achieve recovery point objectives (RPOs) of hours to minutes, depending on the bandwidth and rate that data changes. TrueCopy Extended Distance guarantees that the updates made at primary storage systems are replicated to and written in exactly the same sequence at the remote site.

Figure 4. Hitachi TrueCopy Extended Distance Software



TrueCopy Extended Distance software at the primary site issues the request to create the S-Vol at the remote site. Once it is created the initial copy of all the data from the P-Vol is made. When the initial copy is complete, the relationship is established and maintained until it is split by either being suspended or broken.

TrueCopy Extended Distance uses the Copy-on-Write Snapshot capability to capture the changed data on the P-Vol. On a periodic cyclic basis, a snapshot is taken, the snap is split, and the data in the snap (the differential or changed data) is transmitted to the remote storage system and applied to the S-Vol.



For each cycle, before data is transmitted to the remote site, the primary storage system uses inband communications to issue the request for a new snapshot to be created at the remote site. Each snapshot at the remote site is used to capture the updates to the S-Vol that reflect data from the last split at the P-Vol — in other words, differential data — as it is applied to the S-Vol. The capacity to have multiple snapshots at the remote site provides the ability to fall back to a specific point-in-time copy of the S-Vol at the remote site that reflects a completed transmittal of differential data from the P-Vol — for example, the data as it looked after the last complete transmission. Without the remote snapshot, the data at the remote site would not be usable since the updates within a snapshot from the primary site may not be received in sequence at the remote site. However, when the transmission of each set of differential data is complete, the corresponding snapshot at the remote site along with the S-Vol represents a point-in-time image of the data that is consistent. As a result, the data at the remote site is current as of the last split for which the data in the pool at the primary/local site has been successfully transmitted.

This implementation supports replication of consistency groups — multiple volumes within the same storage system — to the remote storage system. This is possible because the snapshots at the primary site are managed by the local storage system, helping to ensure that the split of the snapshots for all volumes in the group all occur at the same time.

For example, a database application not only updates the actual database records but also causes entries in the database log to enable the changes to be backed out or re-applied in case of a problem. When an update to the database is written by the application, the database system will write a log entry indicating that the update is about to happen. The update is then made to the database and another entry is made to the log, indicating that the update completed. If these logs and databases are on different volumes, and the corresponding updates in separate snapshots, there is the possibility that during replication the split of the snapshot of the volume with the log file indicating the completion of the update could occur before the split of the snapshot reflecting the actual update to the database or the opposite could occur. If the split for the volume with the log file occurred first and was sent to the remote site, there would be no method to guarantee consistency, the differential data for the volume with the log file could be replicated to the remote site, the primary site could have a failure and the snapshot with the update to the database itself may never be received at the remote site. All indications at recovery are that the update completed successfully because the log has both entries for the update even though the update was never actually applied. With TrueCopy Extended Distance software's Consistency Group support, the snapshots for both the volumes with the log files and the database would be split together and transmitted to the remote site and applied. The log and database volumes are consistent and, as a result, any recovery action would proceed correctly.

Because TrueCopy Extended Distance software is storage system-based, it is independent of the host operating system or application and does not require host cycles to operate or to be managed. It is managed with Hitachi tools, which also run on the storage system.

TrueCopy Extended Distance software allows organizations to start small and grow, based on replication needs. From a budgeting perspective, it is easy to expand, as there is a single license per storage system providing for replication of unlimited capacity for the licensed storage system(s). The organization can add storage and replicate more data with no additional charges for the software.

TrueCopy Extended Distance software enables administrators to implement a data protection infrastructure that fosters resilience and supports nondisruptive business operations. It is priced based on the frame, not the amount of storage installed in the frame or that is being replicated, so it has a low cost to scale, and can replicate more data without increasing the license fees. It truly provides remote data consistency with the ability to be back up and running quickly. It is entirely storage systems-based so it is independent of the host processor, the operating system and the application, ensuring that host resources are available for production applications, not consumed by data replication.

Hitachi Protection Manager Software

Hitachi Protection Manager software is the perfect complement to Hitachi replication software for Hitachi modular storage in environments with Microsoft® Exchange and/or Microsoft SQL Server. It provides the ability to automate nondisruptive online backup of Microsoft NTFS, Exchange, and SQL Server environments. It facilitates rapid recovery in the case of data corruption or system outages. Protection Manager software protects the data associated with these mission critical applications by automating disk-to-disk backup and rapid recovery using proven Hitachi replication technologies along with Microsoft Volume Shadow Copy Service (VSS) for Exchange and NTFS and Microsoft Virtual Device Interface (VDI) for SQL Server.

Protection Manager software works with Symantec NetBackup, IBM® Tivoli Storage Manager®, Symantec BackupExec and CommVault®. Protection Manager software can mount and unmount secondary volumes and invoke backup commands. For Exchange and NTFS, the VSS Import server, which also performs verification of the backup, can share the physical server with the backup software's media server.

To integrate the software with existing¹ backup applications, GSS offers a services to deploy Protection Manager software. When implementing Protection Manager software, the configuration should be designed to protect against a local disaster and any failure of the local storage controller. In the case of a local disaster, critical data should be backed up to a remote location. For a local storage controller failure, data should be backed up to another storage controller, either local or remote. Hitachi Storage Replication program products are a prerequisite for the use of Protection Manager software.

ShadowImage software or Copy-on-Write Snapshot software can be used for local replication with Hitachi modular storage. TrueCopy Extended Distance can be used with Protection Manager software for remote replication, but if used with VSS, the configuration and planning must take into consideration that VSS can time out depending on pending in queue or bandwidth.

Protection Manager Software with Microsoft Exchange

For Microsoft Exchange, Protection Manager software provides three backup options:


- Cold — requires the information store to be dismounted, resulting in a service interruption
- Online — implemented using the traditional split mirror backup, which discovers SQL/Exchange and file system objects, maps the relationship from databases of instance to storage logical unit numbers (LUNs) and works with VSS for Exchange and NTFS and with VDI for SQL
- VSS — provides high fidelity backup for Exchange and NTFS²

A variety of options for backup and restore granularity are supported using Protection Manager:

- Individual storage group level, which requires individual storage groups to be assigned per LUN
- Single database level, which requires each pair of database files (.edb and .stm) to be assigned per LUN
- Roll forward
- Point-in-time recovery

¹ Assumes the existing backup application is one of the supported backup applications, which include Symantec NetBackup, IBM Tivoli Storage Manager, Symantec BackupExec, CommVault

² Hitachi Protection Manager only supports Exchange 2007 with VSS. There is no support for cold or online mode for Exchange 2007.



Restoration of a single item can be accomplished through the integration with third party products (for example Ontrack Power Controls). Protection Manager software supports both Microsoft Cluster Server (MSCS) and VERITAS Cluster Server (VCS)³.

Protection Manager Software with Microsoft SQL Server

When used in conjunction with Microsoft SQL Server, Protection Manager software offers the option of creating a full backup through interaction with VDI or creating a differential backup. The backed up copy of the database can be restored to either the original or an alternate instance of SQL. The backup and restore operations can be performed at the individual database level; however, each database must be assigned per LUN.

The combination of Protection Manager software and Microsoft SQL Server supports Extended Maintenance Mode, whereby it will keep all cluster resources online throughout the restore process and limit the impact of database access only to the LUN being restored. For high availability environments, Protection Manager software with SQL Server is supported on both Microsoft Cluster Server and Veritas Cluster Server.

Protection Manager software offers both a command line interface (CLI) and a graphical user interface (GUI). It can discover, map and maintain end-to-end relationships of SQL Objects (instances, databases, file system layouts, volumes, copy groups, LUNs, etc.) with the Hitachi replication copy pairs and provide backup/restore target oriented operations for Microsoft SQL and Microsoft Exchange (Storage Group). It can also send e-mail notifications of the result of copy and/or restore jobs. There is a scheduler that provides basic workflow capability for pre- and post-job scripts.

Benefits

Protection Manager software can deliver significant benefits to the organization, including the ability to significantly shrink backup windows, improve recovery time objectives (RTOs) and increase availability of critical applications to meet service level agreements (SLAs) and reduce unplanned downtime and associated costs. The ability to create backups with little or no impact on production enables more frequent backups, improving RPO. Protection Manager software provides an easy-to-use interface that simplifies implementation and ongoing operations. Hitachi Data Systems Technical Resource Centers, Hitachi, Ltd., and Microsoft support Protection Manager software.

Replication Management

Hitachi Replication Manager software configures, monitors and manages Hitachi replication products for both open systems and mainframe environments. It is a centralized, GUI-based management tool for disaster recovery and in-system replication bundle solutions for both open systems and mainframe environments. It builds on existing Hitachi technology by leveraging the powerful replication capabilities of the storage systems and by combining robust reporting, mirroring and features previously available in separate offerings.

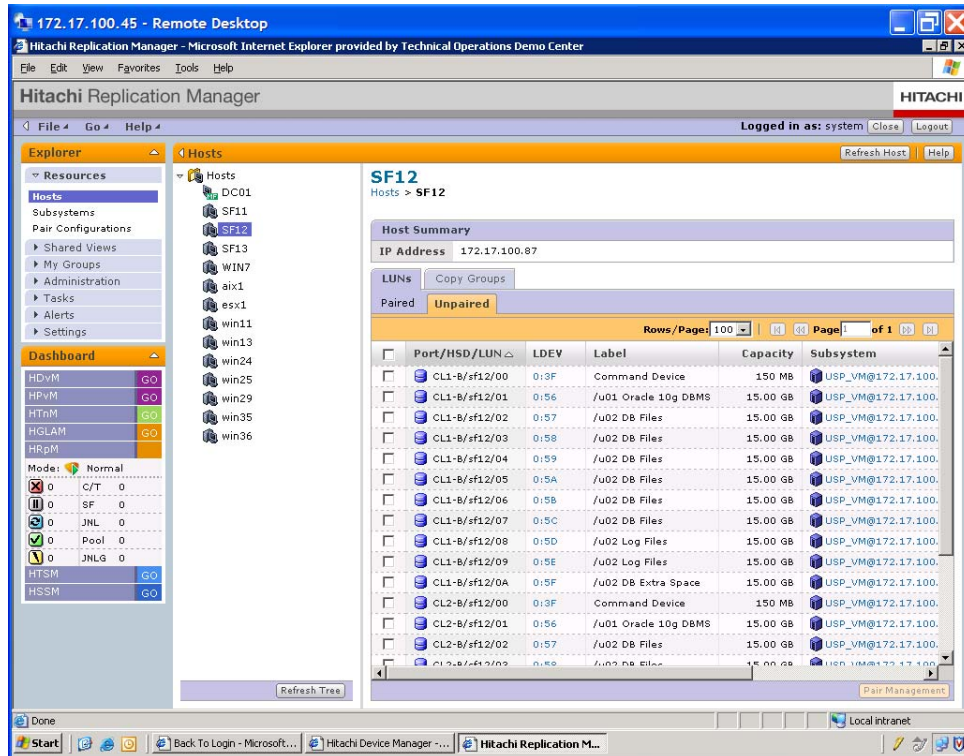
Replication Manager software decreases management complexity while increasing staff productivity. It provides greater control than previously available solutions through a single, consistent user interface — an intuitive GUI manages system, site, region or global replication storage solutions.

With Replication Manager software, replication is no longer restricted to experienced storage administrators. This powerful tool can simplify the management of the replication infrastructure and help achieve RTO and RPO to improve administrative efficiency. It uses a “single pane of glass” for managing both mainframe and open systems, providing mainframe environments with a graphical representation of the replication

³ Veritas Cluster Server with Exchange 2000/2003 is only supported using cold or online modes.

infrastructure. Alerting features, such as e-mail and SNMP traps, allow administrators to monitor the health of the replication environment.

Figure 5. Replication Manager Console Screen Shot



The Replication Manager console, shown in Figure 5, provides access to the configuration and operational information for the replication environment. In the upper left corner of the screen, under “Explorer”, the user can select a variety of views, including resources, administration, tasks, etc. When the user clicks on “Resources” and selects “Hosts,” the right side of the screen provides a list of the hosts and the ability to select each host, and view and manage the replication status of each LUN associated with the host.

All pair operations can be performed from this Replication Manager console. Other installed Hitachi storage management tools can be accessed from the dashboard on the right side below Resources.

Summary

Hitachi Data Systems replication solutions for modular storage provide a wide range of functionality and data protection capabilities. From basic local replication like ShadowImage software and Copy-on-Write Snapshot software to remote solutions like TrueCopy Extended Distance and application specific solutions such as Protection Manager software, Hitachi Data Systems offers the tools to protect critical data assets.

It is important to understand an organization’s specific requirements for protecting data. Not all data is critical. When developing a plan, it is prudent to include different levels of protection for different classes of data. The range of replication offerings for Hitachi modular storage provides the tools to tailor and implement a replication and recovery infrastructure that is right for the specific business.



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 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.emea@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., 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.

All other trademarks, service marks and company names mentioned in this document or Web site are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or limited, 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 Corporation 2008. All Rights Reserved. WHP-317-00 DG December 2008