HDS Storage

HITACHI
DATA SYSTEMS

Backup/Restore Considerations for Multiplatform Storage

by Hubert Yoshida

For computing as critical as your business

Contents

Backup/Restore Considerations for Multiplatform Storage

by Hubert Yoshida

The Business Demand for Backup/Restore is Skyrocketing

Information is no longer a byproduct of day-to-day business activities—it **is** the business. Companies that rely on data and information for survival recognize it as their most important asset. Keeping that information safe is critical to their business success and competitive well-being. Backup/restore is a key component in protecting information. As the amount of information produced skyrockets, and the cost and complexity of backup/restore continues to increase, it becomes even more vital for companies to choose their backup/restore solutions and their enterprise storage carefully.

Not surprisingly, backup/restore is the largest component of the total cost of ownership for storage. The Gartner Group estimates that it accounts for 32 to 54 percent of total costs (see Figure 1). These costs are expected to increase with the need for 24x7x365 operations, the implementation of disaster recovery solutions such as Hitachi Remote Copy, and the trend towards new applications that generate tremendous amounts of data.

Applications such as collaborative computing (Lotus Notes and MS Exchange), data warehousing/decision support, and business process automation (SAP, BAAN, People Soft) dramatically increase the need for backup/restore. Office mail systems are often backed up two or three times a day, and must be recoverable to the document or message level. Business process automation systems are built upon relational databases that provide a single view of data for the entire business from sales and order entry to shop floor control. This puts high value on the ability to rapidly restore data and recover applications.

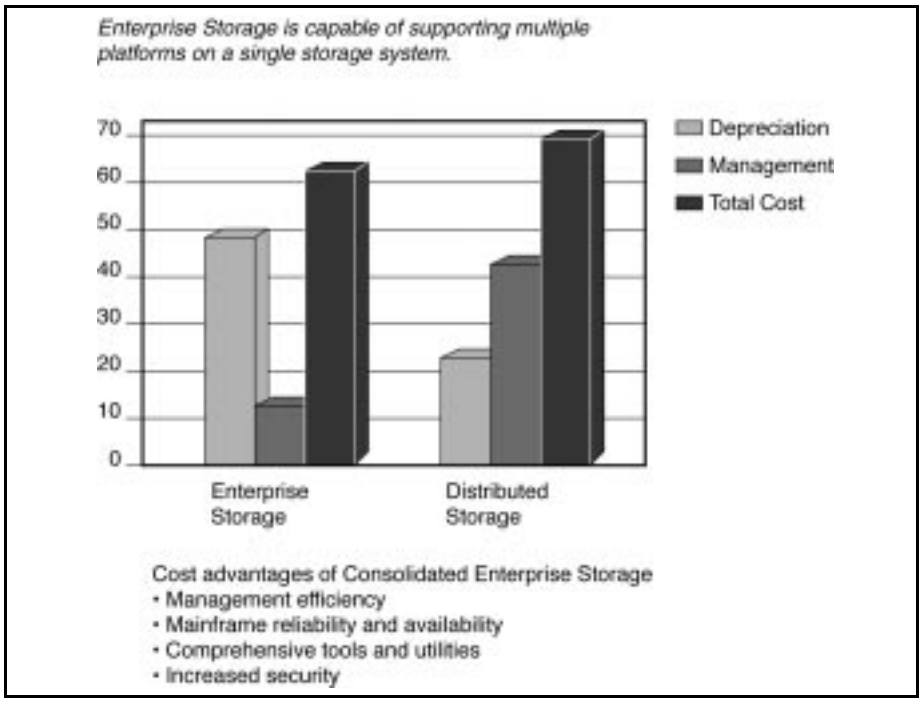

Figure 1: Total Cost of Ownership Model

Hitachi Data Systems Offers Complete Enterprise Storage Backup/Restore Solutions

Hitachi Data Systems addresses the needs of both MVS and LAN customers with backup/restore solutions that are scalable to the enterprise level. Intended respectively for MVS and LAN environments, HARBOR and HYPERtape support a full menu of open systems clients, mainframes, and agents for both on-line and off-line backup of databases. (HDS is a reseller of HARBOR and HYPERtape backup/restore software packages in the U.S. and Canada.) In addition, HDS offers several backup/restore features and services that help reduce the total cost of computing.

HDS backup/restore solutions are designed to meet the following customer requirements:

- Maximum reliability
- Centralized administration
- File-level backup/restore
- On-line backup of popular databases, applications, and messaging services
- Out-of-the-box agents for popular servers such as Oracle Server, BMC SQL-Backtrack, and so on
- Multiplatform backup to a mainframe or open systems server
- High-performance movement of data
- Ability to "snapshot" the backup source in UNIX®/NT as well as MVS
- Ability to simulate or test the recovery scenario.

Hitachi Data Systems is also one of the few vendors who can provide an enterprise storage solution today. In "Demonstrating the Value of Enterprise Storage," IDC's John McArther defines "Enterprise Storage" as consolidated storage that is capable of supporting multiple operating systems and multiple mainframe, UNIX, and NT servers on a single storage system. McArther examines the storage management costs associated with enterprise storage versus distributed storage. Even though

enterprise storage incurs higher hardware and labor costs, its management efficiencies more than offset those costs (see Figure 2).

McArther points out that even if hardware costs for enterprise storage remain twice that of distributed storage, the total cost advantage of enterprise storage is likely to increase over time due to these management efficiencies. Many analysts believe that the hardware costs for enterprise and distributed storage will converge in the near future, which will make the case even stronger for enterprise storage. Factors affecting storage management efficiency include traditional mainframe reliability and availability, automated storage management utilities, tape libraries with tape management systems, asset tracking, capacity planning, performance monitoring, and tuning.

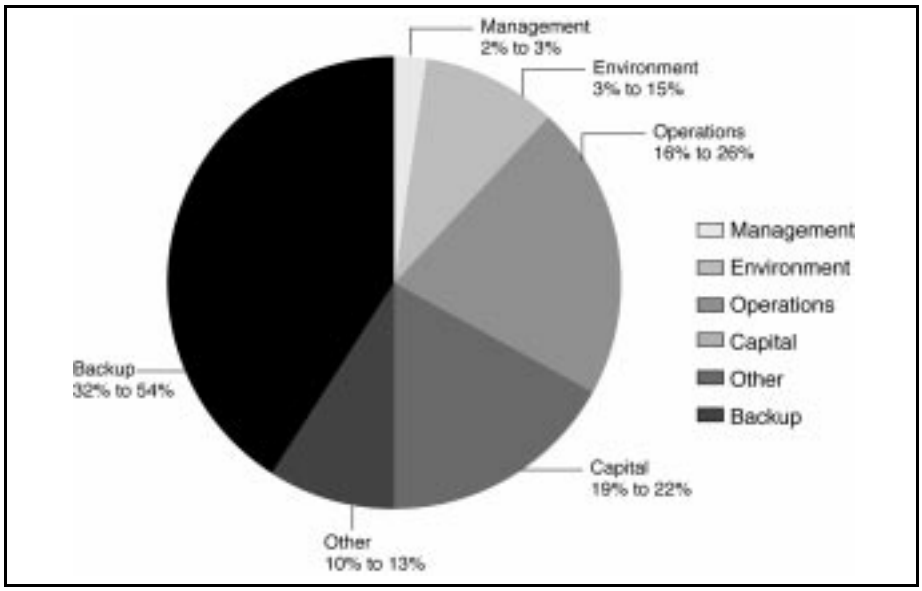

Figure 2: The Value of Enterprise Storage

Other studies have shown that the restore success on enterprise servers is 99 percent, while the restore success on distributed and desktop servers is 25 to 75 percent. McArther concludes that "the justification for Enterprise Storage comes not from simplistic $/MB measures, but rather from improvement in storage management efficiencies, avoidance of costly losses, and enhancement of business value. …Storage buyers and manufacturers that focus solely on hardware $/MB measures and not on overall value will find themselves at a significant competitive disadvantage."

Before we examine Hitachi Data Systems' backup/restore solutions in greater detail, let's take a look at the different levels of complexity, the process of end-to-end recovery, and the constraints surrounding the time available for backup/restore.

Hierarchy of Backup/Restore

The complexity and cost of backup/restore varies, depending on whether it is conducted at the volume, file, database, or application level. The lower levels of backup may be the easiest to execute from a hardware standpoint, but may be the most expensive in terms of the management costs and downtime required to recover the applications which are dependent on that data.

Volume-level backup/restore

This is the easiest to do from a hardware/software viewpoint. However, the customer must implement controls to ensure the logical and temporal consistency of the data on the volume with the application view of the data. In order to have

consistency, a volume-level backup requires that all applications be stopped, and that the data in memory and cache be flushed to the disk. If this volume is part of an application database, it must be coordinated with the backup of the application's logs, indices, repository, updates, exports, and recovery scripts, spread across multiple volumes or generations of volumes. A volume is an MVS concept and datasets are contained within a volume with a description of the dataset recorded in a volume table of contents (VTOC).

In UNIX/NT, the closest representation to a volume is a SCSI address (Logical Unit or LUN) or mount point (C disk, D disk, and so on). Unlike MVS, open systems files will often span multiple LUNs and the file locations are defined by hierarchical directories within the file system. Unless the entire file system is backed up at the same time, the restore of volume-level (LUN) backups is not possible without file-level information from the host file system. Multiplatform volume backups between S/390 and UNIX/NT are available with Hitachi Multiplatform Backup/Restore (HMBR) on the Hitachi Freedom Storage™ 7000 Series. With the HMBR feature enabled, a user can back up any open systems LUN on the Freedom 7000 Series to an MVS or VSE system, using any standard utility such as DFDSS or FDR. HMBR is the fastest way to back up LUNs to a mainframe.

File-level backup/restore

This is more difficult to do since it requires the backup software to know and communicate information on the location and extents of a dataset or file. In MVS, this information resides on the disk in terms of catalogs and VTOCs. In UNIX/NT, the information about files and file structures resides in the host file system in terms of mount points and I-nodes.

File-level backup between platforms requires a communication between a backup agent on a client system that can interrogate the file system and a backup agent on the server that records the file control information for retrieval. As with volume-level backup, additional software and management is required if this file is part of a database or application backup. Solutions for file-level backup between S/390 hosts and UNIX/NT are provided by ADSM, HARBOR, and HYPERtape by Multistream. The file-level transfer of data between UNIX/NT clients with these products is done over a network. HDS and HARBOR have developed a unique agent that provides file-level backup through the Hitachi Freedom 7000 Series. The HARBOR client running on open systems platforms sets up control information (including authorization and file lists) and sends it to the MVS host through the network. When MVS returns acknowledgment, the client sends the corresponding file data to the host through Hitachi Multiplatform Data Exchange (HMDE). This relieves the traffic on the network, and minimizes the backup time by transferring data at channel/bus speeds rather than at network speeds.

Database-level backup/restore

Databases can be hierarchical (like IMS and VSAM) or relational (like DB2, Oracle, Informix, and Sybase). DB2 is the predominant RDBMS on mainframes, while Oracle is predominant on UNIX and NT. Databases consist of logs, indices, data, and scripts. Relational databases can be further defined by tablespace which may contain database tables consisting of columns and rows. Databases that are shared require a locking mechanism to maintain logical consistency.

Standalone or "cold" backup of a database requires that the applications against the database be stopped and all the in-flight data be flushed to the disk before backup is started. The applications must remain stopped until the backup is completed and the database is brought back on-line. In order to reduce this downtime, customers use database utilities to do on-line or "hot" backups. On-line backups still require the applications to be stopped and data to be flushed to the disk. But once this is done to establish a synch point, the applications can resume operations in a read access mode. Some utilities provide transaction logging to allow limited write updates

during the on-line backup. Database vendors are usually the source for these utilities since they require comprehensive knowledge of the database architecture and

application. There are some ISVs (like BMC) who provide products like SQL-Backtrack, which can support on-line backup of multiple vendor database platforms. However, the backup, and, more importantly, the restore of a database requires a database administrator who has a comprehensive understanding of the database and relationships between files. Scripts must be written, tested, and maintained for backing up all the necessary database components.

On-line backup of databases requires speed. The longer it takes to do the backup, the more new transactions accumulate in the logs, and the longer it takes to resynch the database after the backup. HDS offers several solutions to minimize this time. Hitachi ShadowImage can be used to take a snapshot copy for backup of mainframe data. The HARBOR file-level backup agent (mentioned earlier in the file-level backup section) can reduce the backup time for HARBOR on-line backup of databases such as Oracle, Sybase, Informix, and SAP.

Application-level backup/restore

Application backup and recovery involves the backup and recovery of logically related groups of database objects across multiple databases. Backups must be synchronized to generate a consistent recovery point across multiple servers. This requires that all transactions active against target objects be placed on hold, and the objects be switched to read-only mode until the backup is completed. Application backup requires a central management server to control and coordinate backup and recovery. A repository must be maintained on information about all database instances, location of backups, and actions performed against instances. Software like BMC PATROL, PRAXIS OMNIBACK, and CA TNG Framework provide this type of backup.

End-to-End Recovery Process

The recovery process begins with the detection of a failure in the system. If the host system is running with high-availability middleware—such as MC/Service-Guard for HP, VERITAS FirstWatch for Sun, HACMP for AIX,™ or Microsoft Cluster Support (MSCS) Wolfpack for Windows NT—the failure can be detected automatically and the hardware connections for data services automatically switch over to a standby system. Disk arrays with I/O path switch middleware support are essential for this switchover process. The Hitachi Freedom 7000 Series supports high-availability middleware on the major open systems platforms.

Once the hardware has been recovered, the following steps are required:

- If a file system was in use it must be recovered. The time for recovery depends on the file system structure (journaled or hierarchical), the size of the file, number of I-nodes, and any structural changes. Some databases, in particular parallel databases, use raw file mode and therefore have no file system recovery time.
- If a database instance has failed it must be restarted and the configuration files must be read.
- Database recovery begins with a roll forward—reading through the redo log files, and rebuilding and reapplying each applicable transaction. The database must be off-line until the entire redo log files have been read and transactions applied.
- After roll-forward recovery is completed, a roll-back recovery is done to undo any uncommitted or incomplete transactions that belonged to the failed system.
- If disk mirroring was in use, the consistency of the mirrors must be checked. If they are not in synch, they are reconstructed from the logs. This step can be avoided through the use of disk arrays like the Hitachi Freedom 7000 Series.
- Once the database is recovered, the application must go through any of its own recovery. Application recovery depends on the application.

Minimizing the Time for Backup and Restore

The available time to do backups is disappearing as businesses go to 7 x 24 operations and databases grow into terabytes. Customers cannot find the time to do the level of backup that they need.

Once a week versus once a day

It is faster for an application to recover a database that is backed up daily versus one that is backed up once a week. Some customers try to reduce the time spent on backups by doing a full backup once a week, along with daily incremental backups. The frequency of backups is a tradeoff against the likelihood of a failure and the time available to do a recovery.

Restore versus recover time

Recovery requires more than just the restore of a volume or dataset/file. It requires a detailed understanding of the nature of the failure, and a determination of what data needs to be recovered, where the backup resides, and the structure of the database or application. It is possible to have all the pieces available for recovery and still not have the knowledge to put them together correctly. An automated recovery process could ensure consistency and completeness in the backup/restore, which would reduce recovery times. Software that goes beyond volume-level or file-level backup/restore and supports application-level backup/restore can help to reduce the time to recover. The installation of a backup/restore solution must be properly planned to ensure complete and timely recovery. HDS Storage Management Services can assist customers with this process.

Snapshot copy backup

On-line backups are the least intrusive, which is why they are extremely popular for databases that require 24-hour availability. However, because the DBMS must log transactions during the backup, performance is degraded both during the backup and after the backup, when the logged transactions are applied to the actual data. In order to minimize the time to back up a volume, customers now have the ability to "snapshot" (take a point-in-time copy of) a volume and do the backup from the snapshot volume, freeing the original volume for production use much sooner.

The Hitachi Remote Copy feature allows users to create a primary and secondary copy pair. This copy pair can be suspended and deleted, and the secondary volume can then be used as a snapshot copy. Updates continue on the primary volume while the secondary volume is being backed up. Both primary and secondary volumes continue to be RAID-protected. Products like FDR InstantBackup can be used to create an off-line backup of the secondary volume. After the backup is completed, the secondary volume is brought back in synch with the primary LVI.

HDS will provide a snapshot copy capability for open systems volumes with its Hitachi Remote Copy/Open feature in the second half of 1998.

Parallel database servers

For some applications, a single parallel database can provide better performance and faster recovery than a non-parallel database. Oracle Parallel Server, Informix XPS, Sybase MPP, and DB2 PE are examples of parallel database servers. Parallel database servers access one large database from multiple hosts that are loosely coupled together through host clustering middleware such as MC/LockManager for HP-UX, Ultra Enterprise Cluster PDB for Sun, and HACMP for AIX. A non-parallel database must restart the database and the applications after a host failure. In a parallel database configuration the application and database are able to remain running, so that a

restart is not needed. Also, since the transaction load is distributed among the servers in the parallel cluster, the transaction log for any one system is smaller and will be recovered faster. Support for parallel database servers is planned for the Hitachi Freedom 7000 Series.

Hitachi Data Systems—For Computing as Critical as Your Business

HDS, in collaboration with industry-leading software providers, strives to address customers' backup/restore requirements with solutions that reduce the total cost of computing. In addition to the HARBOR and HYPERtape backup/restore software packages, HDS offers several features and services to achieve this goal. Hitachi ShadowImage and Hitachi Remote Copy can be used for the creation of a snapshot copy for near-instantaneous backup. Hitachi Multiplatform Backup/Restore delivers high-performance volume-level backup of open systems data. And HDS Storage Management Services consolidate the methodologies needed to optimize storage resources.

Global access, availability across multiple platforms, centralized management, information protection, rapid response to change—these IT capabilities drive the design and development of Hitachi Freedom Storage. Hitachi Data Systems delivers the comprehensive, customizable, manageable, application-enabling hardware and software to help companies align their IT needs with their business requirements. Building and protecting the mission-critical environments of the world's largest corporations is Hitachi Data Systems' primary strength.

Hitachi Data Systems

www.hds.com

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia-Pacific Headquarters
11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325 3300
info@hds.com.au

Canada Headquarters
380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdscanada.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-61-8000
info@hds.co.uk

Latin America Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
lad@hds.com

U.S. Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1066
ussales@hds.com