

Storage Management Solutions for Distributed Computing Environments

Strategies for dealing with the vast amounts of data generated by today's information technology environments involve more than just larger and larger disk drives. They include the right combination of different storage devices to deal with offline, nearline, and online data storage and scalable management software.

by Reiner Lomb, Kelly A. Emo, and Roy M. VanDoorn

Storage management is fast becoming one of the most important issues information technology (IT) managers face today. With data accumulating at enormous rates, and with end users demanding faster access to more information, storage management has moved from an operation that was done only at night to a mission-critical concern that requires full-time attention.

Storage management consists of all the activities related to the effective deployment, accessibility, and use of stored information across a computing infrastructure. Storage management involves several major disciplines, including backing up and restoring data, storing data online across multiple classes of storage devices such as disks and tape, archiving data for legal and historical purposes, and managing storage resources such as tape or optical media for optimal use. Managing these storage disciplines takes an effective combination of organization, processes, and technologies to meet end-user data-availability expectations.

In today's distributed computing environments, IT managers need consistent storage management strategies and processes across the enterprise. In addition, storage management processes cannot be separated from an integrated network and system strategy. Therefore, IT managers need complete solutions that integrate the various storage management components and technologies such as databases, file systems, storage peripherals, storage management applications, and network and system management strategies.

In the past, storage management solutions have been proprietary (mainframes) or piecemeal (early client/server point products), with specific peripherals working only with specific software and hardware. It was difficult to expand a solution to meet the demands of rapidly growing collections of data.

In this article we will describe trends driving storage management technology and the components that make up an ideal storage management solution. Finally, we'll introduce HP hardware and software products, services, and partners and describe how they work together providing storage solutions for our customers.

Storage Management Trends

Traditionally, the task of storage management was done after work hours when the system could be brought down for storage management functions such as backup and archiving. Today, much more data is generated, and storage management solutions need to provide much greater data availability and reliability. Complicating storage management are variables that determine data throughput and access. These variables include disk capacity, CPU, input/output channels, device speed, networks, and software (Fig. 1). New ways of transferring and storing large amounts of data without downtime have to be developed.

Probably the most important driving factor in storage management today is that customers demand continuous accessibility to huge amounts of data, very often in the terabyte range. You can see this demand occurring in the increased use of the Internet and online services such as CompuServe and America Online, and in the emergence of new applications such as imaging and multimedia. In the past, data accessibility was a fairly simple process when mainframes were the primary storage devices and the only limitation was disk size. Today, the answers to storage problems cannot be provided simply by installing a bigger disk on a central server.

As customers reengineer their businesses, many are choosing to migrate away from the mainframe via "mainframe downsizing." Mission-critical applications are moving to open systems, and the management of client/server workgroups is being consolidated across LANs and WANs. An enormous amount of company-sensitive data, which used to be under central control and located in the data center, is now distributed and available on the network (Fig. 2). Published market numbers show that the average amount of distributed data has surpassed the average amount of data in the data center. Companies must begin viewing storage management as integral to their network and system management solutions.

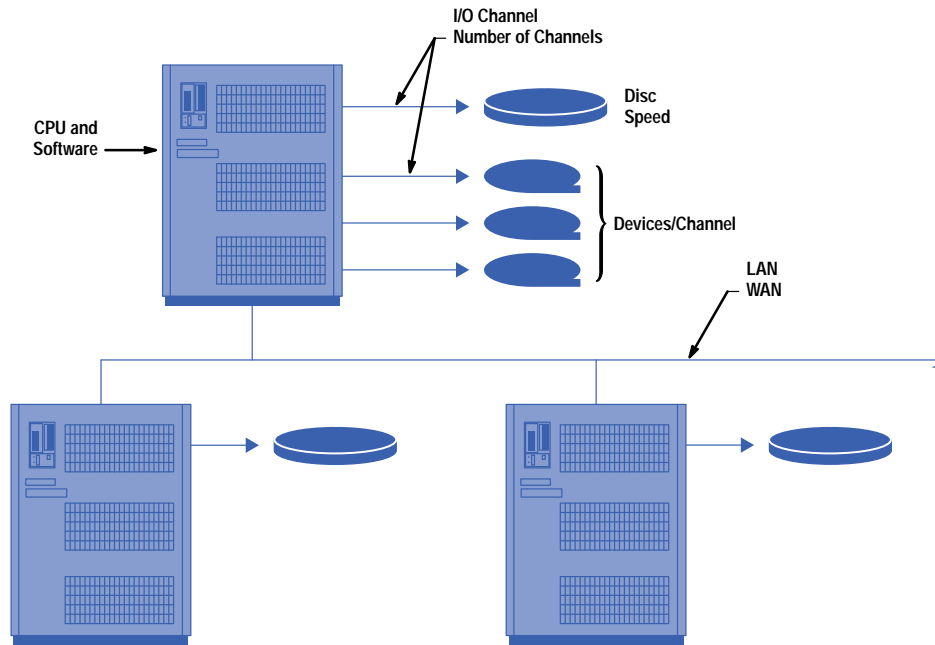


Fig. 1. Components in the chain that have an impact on data throughput.




Fig. 2. Decentralized storage management committed to the same service level as a mainframe environment.

In addition to all the challenges raised by managing storage on distributed systems, IT managers must deal with the reality that the amount of data being stored is outstripping the network's capacity to handle it efficiently (Fig. 3). For example, a company might need to back up 100 Gbytes of data in an hour. As the storage staff looks for solutions, they see processor performance improving faster than disk performance. They also see the performance of both disks and processors outstripping the performance of the installed network infrastructure. At the rate network infrastructure is improving, it will be a huge challenge to catch up to processor performance.




Fig. 3. *The volume of data being stored is outstripping the ability of network services to deal with it.*

If IT managers try to win this contest with only the traditional approach of centrally stored data, they will lose the storage management race. Instead, today's storage management solutions must allow distributed storage management to be performed centrally, decentrally, or in a hybrid fashion, depending on a company's policies and needs.

To solve the growing issue of storage management, they need to understand what constitutes a storage management solution.

Storage Management Requirements

The ideal storage management solution, which is made up of complementary software, systems, and peripherals, is integrated, scalable, and modular, and allows the solution to be implemented in phases and expanded over time. A flexible solution addresses both mission-critical enterprise-wide requirements and business-critical desktop needs. At the same time, this solution must be easy to use and robust, and must provide quick, reliable access to data.

The fundamental requirement of any storage management system is to provide data accessibility to all users, regardless of where and how the data is stored. To make data quickly accessible, yet store it efficiently, customers need a complete, integrated set of storage management functions, including backup and recovery, archiving and retrieving, hierarchical storage management, and media management.

In addition, an enterprise-wide storage solution must allow various storage management applications and peripherals to manipulate and share media in a consistent manner. It must also provide an easy and standardized way to access the various storage devices, library systems, and silos. Many companies are dedicating servers to specific tasks such as backup and restore servers or archival and retrieval servers. A storage solution must be optimized so data is stored and moved in the most efficient manner.

Other more generic services for storage management include a central policy definition and a single point of control.* Lights-out operation and unattended remote backup are also key to many storage management solutions.** Storage management solutions are most manageable when integrated into management systems such as HP OpenView, which provides integrated network and system management services dealing with monitoring, problem management, and configuration and change services.

Backup and Recovery

One of the most important needs in enterprise-wide storage is backup and recovery. Very early in a solution deployment, IT managers must establish a backup and recovery policy that provides the appropriate level of data integrity. This policy must ensure that critical data can be completely and quickly recovered from a backup even in the event of a disaster. Equally critical is minimizing planned downtime, or completely avoiding downtime, to create a backup while keeping user applications up and running.

* A central policy describes a set of features that allow an administrator to define policies about how distributed storage is to be managed from a central management console. For example, an administrator defines for a networked environment which data needs to be backed up, when it will be backed up, which device will be used for backup, and so on.

** In the IT community lights-out operation means that an IT environment can run without local operators or administrators.

Archiving and Retrieving

The main reason for implementing archival and retrieval solutions is the need to keep data long-term and guarantee retrieval when access is required. The data is typically copied onto a different medium such as tape or optical disk, while the original copy is deleted from magnetic disk. Archived data is not frequently accessed, but sophisticated retrieval mechanisms need to be available. In many cases, archiving data is required for legal or internal auditing purposes. The archiving procedure includes storing data that logically belongs together in long-term storage, such as a finished project, a finished design, or a client record.

Hierarchical Storage Management

Hierarchical storage management, or HSM, efficiently manages data stored on magnetic disks, optical disks, and tapes. Depending on cost versus performance requirements, data is kept on one or more of the different storage hierarchy levels and migrated transparently among the storage media according to customer-defined policies. An HSM system reduces ongoing storage configuration tasks, such as moving data manually between levels in the hierarchy and subsequent management costs. It also eliminates frequent storage maintenance, such as manually archiving files onto tape to free disk space, and it helps reduce the need to acquire more expensive media, such as magnetic disks, for infrequently accessed data. For example, files are migrated from disk to tape or optical storage if they are not accessed for a certain time period. Statistical data about access patterns can help to define the right migration policy. Also, keeping statistical data about migration patterns and creating appropriate reports will help to implement the right storage management policies for an organization.

Media Management

The storage services discussed above handle copying or moving data onto media or retrieving data from media. Media management, which keeps track of removable media such as tapes or optical devices, deals with the medium itself and not with the data on the medium. A media management system protects data on the media and makes the media pools available to storage management applications. Typical media management functions include mount and unmount media, rotate media, and provide statistical information about the media.

Most of today's backup, retrieval, archival, or HSM products have their own integrated media management functionality dedicated to a specific product. Enterprise storage management solutions require generic media management services delivered in an integrated way, so that media use can be managed and optimized across applications and systems.

Enterprise-Wide Storage Management

IT departments also require consistent and effective management capabilities for storage management across the enterprise environment. To provide these management services, a complete storage solution must provide:

- A single point of control, which is consolidated console management, including:
 - Central policy definition
 - Central monitoring and problem management
 - Central configuration of storage
- Multivendor availability and support
- Scalable, modular services
- Integration with an industry-standard network and system management framework
- High availability of key storage management components.

HP Storage Management Solutions

HP can offer many different solutions to an organization's storage needs because of the combined effort of major HP business organizations in the areas of network and system management, storage peripherals, and UNIX[®] servers.

However, each customer's needs for storage management solutions are different. No one vendor can provide a single solution for every environment. Rather than create a monolithic, proprietary solution, HP is working closely with third-party partners and many diverse HP divisions to create an open, standardized environment in which many vendors can participate in creating solutions.

HP Storage Management Software

Central to HP's storage management offering is the software that links all the other pieces of a storage management solution together. HP's two leading storage management products, HP OpenView OmniBack II and HP OpenView Omni-Storage, are client/server solutions that give IT managers the flexibility to manage distributed storage centrally and delegate management responsibility to distributed sites or departments. HP OmniBack II products address issues associated with data loss and protection, and HP OmniStorage products address issues associated with space management.

HP OpenView OmniBack II. HP offers two backup solutions: HP OpenView OmniBack II for Workgroups, which provides an entry-level backup solution ideally suited for the workgroup environment, and HP OpenView OmniBack II, which provides a comprehensive backup management solution to cover all sizes of environments, including the whole enterprise.

The HP OmniBack II architecture (Fig. 4) consists of three major pieces:

- Backup Manager. This module centrally administers and controls the backup environment.
- Backup Device Servers. These servers run on the system to which the backup device is connected. A backup environment can have many backup device servers.
- Clients. All systems being backed up need a client to invoke the backup utilities.

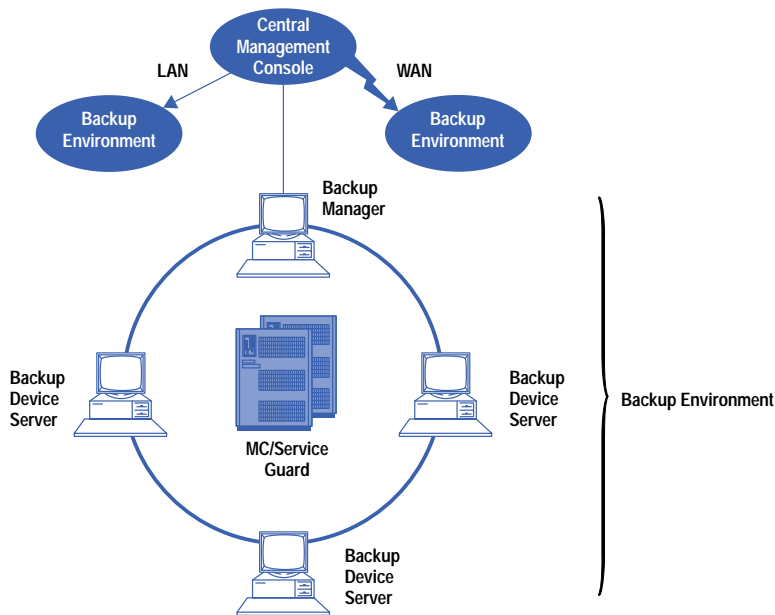


Fig. 4. The backup environment and components for HP OpenView OmniBack II.

All three components can run on the same system or can be distributed. OmniBack II has its own management interface and can be run inside or outside the HP Openview management interface.

OmniBack II is a scalable and flexible solution. Through its policy-driven, centrally managed, automated backup capabilities, OmniBack II reliably protects data distributed throughout the entire network. Easy-to-use backup and restore functionality provides management for desktop PCs to UNIX-based business servers. In combination with HP OpenView IT/Operations, administration and problem management for the entire enterprise can be centralized.

Sophisticated media and device management combined with support for mainframe-class library systems, including silos, make OmniBack II the ideal solution for data centers.*

Increased uptime for application and database servers can be achieved through high-performance offline backup, requiring only a minimum of application downtime. This can be extended to 100% application availability through online backup of business data.

The main features provided by OmniBack II include:

- Network backup and recovery
- Support for a broad range of devices and libraries
- Online backup of applications and databases such as SAP/R3, Oracle, and Sybase
- Sophisticated media management
- Support for major UNIX and PC platforms, including Windows NT
- High-performance backup and recovery from multiple drives in parallel, each running at its full native throughput
- Integration with HP OpenView IT/Operations
- Integration with HP OpenView OmniStorage, HP's hierarchical storage management solution.

HP OpenView OmniBack II for Workgroups. OmniBack II for Workgroups is a complete solution that offers everything needed for a low-administration, automatic, unattended, and reliable network backup and recovery solution. It is targeted toward

* Silos are very large tape libraries.

smaller multivendor computing environments without dedicated administrators. OmniBack II for Workgroups includes the following features:

- Automated and reliable network file system backup and recovery
- Sophisticated and automated media management, autoloader support
- Support for all major UNIX and PC platforms
- Easy-to-use intuitive graphical user interface with many built-in browsers and selection lists.

HP OpenView OmniStorage. OmniStorage is HP's hierarchical storage management solution. It offers benefits in environments where a significant amount of data needs to be online, but where not all of the data is frequently accessed.

OmniStorage provides high-capacity, cost-effective online storage by supporting HP's broad range of optical libraries and the newest tape libraries. According to policies defined by the customer, files are automatically and transparently migrated among the levels of storage hierarchy.

Fig. 5 shows a typical environment in which OmniStorage runs. The two main pieces of OmniStorage are the manager and the clients. The manager administers and controls the storage environment, and the clients invoke the OmniStorage functions on behalf of users.

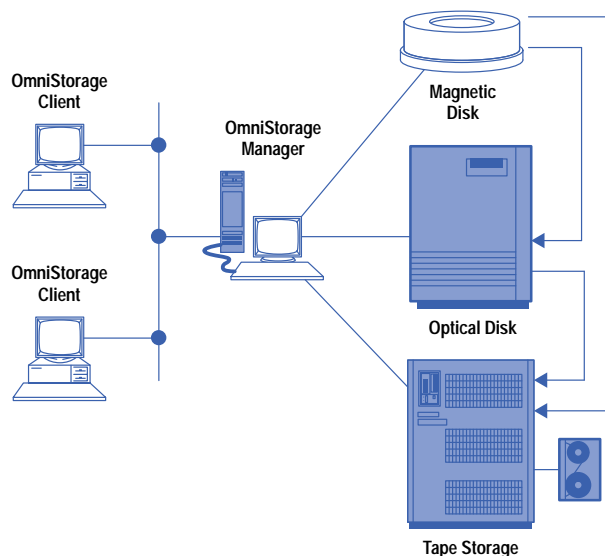


Fig. 5. The components and environment for HP OpenView Omni-Storage.

OmniStorage is tightly integrated with HP OpenView IT/Operations, providing easy administration and problem management of multiple OmniStorage installations from a central workstation console. OmniStorage also integrates with OmniBack II for automated backup and recovery of the HSM environment. However, OmniStorage can run as a standalone product, which allows customers to implement storage management in phases.

OmniStorage provides optimal performance if users frequently access only a subset of the data. Additionally, Omni-Storage can be used for databases if they are based on a file system and if major parts of the database, such as decision support systems, are not frequently accessed.

Finally, OmniStorage provides the following features:

- Policy-driven automatic and transparent file migration
- Network migration for HP-UX* and Solaris operating systems
- Additional multivendor support through NFS
- Exceptionally fast rebuild capabilities in case of data loss
- Configurable demigration strategy
- Archival to WORM (write once, read many) disks
- Integration with HP OpenView OmniBack II
- Integration with HP IT/Operations
- Support for data warehouse environments.

HP OpenView Solutions

HP OpenView's solutions are part of a strategy for managing multivendor networks, systems, applications, and databases from the mainframe to the desktop PC. The HP OpenView portfolio and companies that provide network and system

management solutions (solution partners) give IT managers the tools to control and manage all enterprise resources and devices centrally, while reducing the cost of systems operations and administration. Besides OmniBack and OmniStorage, more than 250 HP OpenView-based management solutions from HP and solution partners integrate with a complete set of common management services to help customers improve service and reduce operation costs.

Central to the HP OpenView products is a user interface that provides a focal point from which the IT staff can manage computer systems and network devices. Although control is centralized through the interface, management functions can be distributed across the enterprise. More important, flexible, distributed interfaces allow several operators and administrators to be involved in the process of IT management.

As an important part of the HP OpenView solutions framework, HP OpenView IT/Operations provides centralized operations and problem management with distributed intelligence across multivendor platforms. With intelligent agents (managed nodes) installed throughout the enterprise, IT/Operations collects up-to-date, accurate information to provide 24 × 7 (24 hours a day, seven days a week) uptime for mission-critical applications. IT/Operations-managed nodes gather information, messages, and monitoring values from a variety of sources. Filters and thresholds ensure that only relevant information is forwarded to the central management system and presented to the responsible IT/Operations operators.

HP and Third-Party Storage Peripherals

The HP 9000 supports mass storage products that provide online, nearline, and offline storage capabilities. The primary differentiator among these three categories of storage is access time. A storage device is considered online when the data access time is a fraction of a second. Nearline storage devices usually access data in the range of a few seconds to a few minutes. Offline storage devices typically require many minutes to hours to access data. Some offline storage strategies that require retrieval from a storage vault may take days before the data is available to the user.

HP and its partners can provide a wide variety of products to meet the individual needs of specific customer environments. These products can be mixed and combined with HP's storage management software to provide the needed end-user solutions.

Online Storage

HP offers two classes of online mass storage products: single-spindle disks and disk arrays.

Single-Spindle Disks. Single-spindle disks* offered by HP are either embedded in the host systems or provided externally within storage enclosures. These disks provide high-capacity, nonvolatile, fast-access mass storage. Single-spindle drives operate at 7200 rpm and are currently available in capacities of 1.05 Gbytes, 2.1 Gbytes, and 4.3 Gbytes as fast/wide differential drives.

These new drives are available as embedded devices in all of the HP 9000 servers, more than doubling the internal online storage capacity. With the addition of the 4.3-Gbyte drive, 21.5 Gbytes of external online storage capacity can now be housed in a single enclosure rack with up to 160 Gbytes in a 1.6-meter-high cabinet.

HP External Storage Enclosures. HP offers two families of storage enclosures for online storage: the HP 6000 SCSI mass storage family and the HP high-availability storage system. HP's high-availability storage system is based on a package design that delivers flexibility and ease of use while providing critical functionality to meet the needs of the enterprise. The system provides excellent availability, hot-pluggable power supplies, dual power cords, cooling fans, and hot-pluggable storage modules. The subsystem connects to the server via dual SCSI buses, increasing reliability and enabling disk mirroring in the same enclosure.

HP Disk Arrays. A disk array is a storage system consisting of multiple disk drive mechanisms under the command of an array controller that communicates with the host (Fig. 6).¹ The key benefit of disk arrays is high data protection. Arrays also provide high storage capacities, connectivity, and configuration flexibility.

HP currently offers three primary disk array families associated with the HP 9000 business server product line. The first family is the HP high-availability disk arrays Model 10 and Model 20, which have a raw capacity of 6 to 80 Gbytes, support RAID levels 1 and 5, and have dual and hot-swappable controllers and redundant cooling and power.**

The second disk array offering is EMC's Symmetrix 3000. The Symmetrix 3000 is a high-performance integrated-cache disk array designed for online storage. As such, the Symmetrix 3000 provides a high level of online performance, an online capacity of up to 1.1 terabyte, and manageability and high availability to HP 9000 business servers. The result is a mainframe-class data storage solution that is simple to manage and is delivered in a high-performing, scalable, protected, and open architecture.

The final disk array is the fault-tolerant, self-configuring, high-performance HP disk array product with AutoRAID technology. The HP AutoRAID disk array eliminates the need for system administrators to understand RAID levels.

* A collection of disk platters on a single spindle.

** RAID = Redundant Array of Independent Disks.

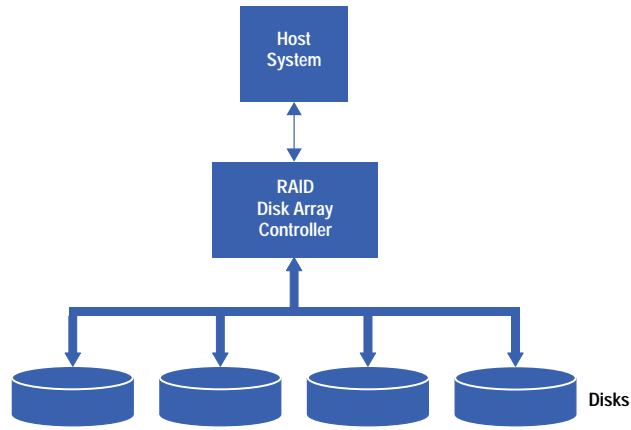


Fig. 6. A typical RAID architecture.

It dynamically adapts to the system's workload, thus optimizing for performance and cost. Finally, it offers a raw capacity of up to 24 gigabytes.

Nearline Storage

Most files and applications stored on hard disks are never used. Thus, the major benefit of hard disks—high-performance access—is squandered on dormant data. Cost-sensitive environments would be better served by a hierarchical storage management solution in which active data is stored on hard disks, while dormant or infrequently used data is cost-effectively stored offline or nearline in media such as optical disks.

HP SureStore Optical Storage Products. HP offers a broad family of optical disk drives, ranging in capacity from 40 Gbytes to 618 Gbytes. HP offers multifunction magneto-optical drives with rewritable and WORM disks. A rewritable optical disk can be written up to 10 million times. A WORM disk can be written once but cannot be erased or overwritten, adding a higher degree of security.

The main features of these nearline storage devices include:

- Fast, near-hard-disk transfer and seek times
- High capacity
- Low risk—no disk crashes with optical disks
- Online data availability on a random-access device
- Online drive replacement—provides assurance that the optical system is persistently available
- Removable media
- Long life—provides more than 100 years of media life without maintenance.

Offline Storage

HP's range of offline storage products provide high speeds and large capacities to meet the increasing demands of HP's high-performance workstations, network servers, and multi-user systems. HP has combined its reliable DAT products with an industry-leading autoloader design and networking software to give customers the flexibility they need for complete automated network backup.²

HP DAT Products. HP offers the latest DDS-2 tape drives in addition to DDS and DDS DC drives. The new DDS-2 format, combined with 120-meter tapes, has a native mode capacity of 4 Gbytes. With data compression, customers can typically store 8 Gbytes on a single tape. For unattended or lights-out operation, a six-cartridge autochanger is available to rotate media for full and incremental backup and restore operations. HP also supports 8 mm and QIC tape drives. The main features of these offline storage products include:

- Unattended backup
- High capacity with high reliability
- Easy storage in a fireproof safe according to industry standards
- DDS format can be interchanged with different manufacturers' tape drives.

Digital Linear Tape. HP 9000 servers support the HP DLT (digital linear tape) Library 4/48. This library consists of four Quantum DLT/4000 drives, accommodating 48 20-Gbyte tape cartridges and providing greater cartridge capacity than the DDS format. The DLT/4000 is a 0.5-in cartridge streaming tape with a capacity of 40 Gbytes per cartridge (with 2:1 compression), and a sustained transfer rate of 3 Mbytes/s. The HP DLT Library 4/48 enables fast, unattended backup of over

100 Gbytes of data within the brief windows of time available for backup in high-end OLTP (online transaction processing) and decision support system environments (Fig. 7).




Fig. 7. System capacities of different offline backup devices based on the backup window (i.e., the amount of time available for backup).

The main features of digital linear tape include:

- A native-mode data transfer rate three times faster than competing technologies
- Greater media and drive head longevity
- Sophisticated tape indexing for fast-streaming file searches and restoration
- Higher compression ratio for most data types.
- Driver support provided by HP for the 18-track StorageTek tape drives (model 4781) and the 36-track single-ended tape drives (model 4791).

3480/3490 Compatible Tape Subsystems, Libraries, and Silos. HP provides driver support for 18-track StorageTek tape drives (Model 4781) and 36-track single-ended tape drives (Model 4791). Additionally, StorageTek offers Timberline 9490, a fast wide implementation of the Model 4791 with a 6-Mbyte/s drive. These drives are compatible with all StorageTek silos. HP supports StorageTek silos, including one with a 500-cartridge capacity and 90 cartridges/hour (upgradable to 1000 cartridges and 350 cartridges/hour) and another with a 6000-cartridge capacity and 350 cartridges/hour. Both connect to other devices and silos for easy growth.

HP 9000 Business Servers

Today's open systems for critical business computing environments require three essential elements. First, they must provide the storage management, data integrity, security, and manageability that information technology managers have come to expect in running business-critical applications on centralized processing systems. Second, they must provide connectivity and compatibility with the growing base of PC desktop users. Finally, they must offer flexibility, performance scalability, and technical innovation to keep up with emerging application demands.

As the leading open systems platform, HP 9000 business servers offer the benefits of all three elements in a single, unified, UNIX-based platform. The HP 9000 server platform is able to support environments of all sizes, ranging from workgroups and replicated sites to the departments and data centers of large enterprises. For storage management, the HP 9000 business servers offer the following features:

- Highly available and reliable systems environment
- Excellent data-movement management
- A dedicated storage server architecture that is designed for optimal database and file management
- Scalable from desktop to data center
- Hundreds of partners that ensure customizable solutions.

Storage Solutions for the Enterprise

To help understand how to plan a complete storage management solution, we have looked at scenarios common to many companies struggling with storage management demands. In each of the following examples, we'll examine the needs specific to each environment and the needs of centralized storage management, starting with small workgroups, and building up to the enterprise level. Then we'll show how HP and its partners can provide a unified solution.

Independent Workgroups

Many workgroups implement their own backup and recovery solutions. These solutions are typically managed by a part-time administrator. Major requirements for backup and recovery solutions include ease of use and automation. OmniBack II for Workgroups is the best backup and recovery product for independent workgroups. Combined with HP's low-cost, high-performance business servers and HP's DDS II device libraries, backup procedures can be automated and centrally controlled within the workgroup. OmniBack II for Workgroups provides easy and fast restoration of files and the potential to expand the workgroup or even consolidate multiple workgroups.

Solution Elements. The storage management solution from HP for independent workgroups includes:

- HP OmniBack II for Workgroups, with easy-to-use backup and recovery software for small environments
- HP 9000 Class D and E business servers for backup
- HP's DDS II autoloader as a cost-effective tape library.

Distributed Client/Server Workgroups

As information technology departments consolidate workgroup management, they need more centralized storage management for distributed, heterogeneous workgroups. Two main objectives of this scenario are to increase end user productivity by providing homogeneous and powerful storage services, and to increase operator productivity through central control and administration of those storage services over the LAN. Significant savings can be achieved through intelligent resource sharing and reduction of operational overhead.

OmniBack II centrally manages the complete backup and recovery process of large numbers of distributed workgroups by dividing large numbers of backup nodes into multiple manageable backup domains. Central control can be maintained at the enterprise console while delegating backup and recovery tasks to the individual end-user departments. OmniBack II can automate the complete backup process of distributed client/server workgroups.

In addition, data on shared file servers and on client disks grows dramatically. Migrating infrequently accessed data onto different storage media such as optical disks or tapes becomes an administrative nightmare. OmniStorage helps to increase the online storage capacity of clients and servers while keeping storage administration costs under control.

Solution Elements. The storage management solution from HP for distributed client/server workgroups includes:

- Centrally controlled backup and recovery for heterogeneous workgroups with OmniBack II
- Automated backup based on HP's DDS II Autoloader or DLT libraries
- HP 9000 business servers used as reliable and high-performing backup, restore, and HSM servers.
- Unlimited online storage based on OmniStorage combined with HP's optical or tape libraries
- Sophisticated problem management with HP IT/Operations.

Regional Distributed Systems and WAN Connections

Companies with branch offices in the retail or financial industries, for example, often have regional distributed systems connected via WANs. IT departments for these companies need to run the IT infrastructure of the branch offices without operators or administrators. Remote control and administration of storage services over the WAN are essential.

OmniBack II defines each remote branch office as a backup domain with one or more local backup servers. The complete backup and recovery administration process and control can be performed from a central console via WAN. By choosing the appropriate devices, with sufficient capacity for each of the remote offices, the backup and recovery process for the branch offices can be performed remotely.

Solution Elements. For regionally distributed systems the storage management solution from HP would include:

- Central backup and recovery for remote sites with OmniBack II
- HP 9000 Class D and E business servers for reliable backup and recovery at branches
- HP's DDS II autoloader for automated tape handling
- HP IT/Operations integration of OmniBack II for sophisticated central problem management.

Data Center and Mainframe Downsizing

When customers migrate from the mainframe to open systems, they expect the same functionality and scalability in storage services as they had with the mainframe. The IT department is expected to continue to provide the same level of assurance that computer services are available, reliable, and secure.

For example, a major multinational company using the HP-UX operating system with large SAP/R3 projects based on Oracle requires efficient, reliable, unattended automated backups and restores. These backups are in the multiple terabyte-per-week range and will move into 24 × 7 (24 hours a day, 7 days a week) operation. OmniBack II is a key part of the solution because it gives the customer online backup by integrating with either the SAP/R3 online utility or the Oracle database utility. OmniBack II also integrates with IT/Operations. This customer has complete centralized management of all backup sessions and devices with the option of managing a partially or fully distributed environment.

Because of OmniBack II's modularity and integration with other online backup services, such as those provided by Oracle and Sybase, customers can configure OmniBack II to expand to match their growing infrastructures. Many customers also may choose to add HP's MC/ServiceGuard to ensure high availability of stored data.

Solutions Elements

Data centers needing high-end backup and restore are presented with the following storage management solutions from HP:

- High degree of automation with OmniBack II and StorageTek's Tape Silo
- High backup and restore performance using StorageTek's Timberline tape drives
- High reliability through HP 9000 systems running the HP-UX operating system
- SAP/R3 online backup with OmniBack II integration
- Enterprise monitoring and problem management through HP IT/Operations integrated with OmniBack II
- Full consulting, from investigation through implementation, by HP's professional consulting services
- HP MC/ServiceGuard.

Data Warehousing: Scalable Storage Infrastructure

Much of a company's data remains valuable but does not need to be online and available all the time. Typical storage capacity requirements of data warehouses range from tens of gigabytes to several terabytes. Storage managers need the ability to move this data cleanly from primary to secondary storage and back to primary temporarily, as needed.

The combination of HP's business servers, magnetic disks, optical disks, and OmniStorage offer an ideal storage infrastructure for data warehouse environments. This solution offers efficient storage hardware costs and high scalability for large storage capacities.

The ideal ratio between magnetic and optical capacity depends on environment. Ratios in the range of 1:5 to 1:10 (i.e., 1 Gbyte of magnetic disk capacity assigned to 5 or 10 Gbytes of optical capacity) have been implemented successfully.

Solution Elements. HP provides the following storage management solutions for data warehouses:

- Automated data migration with OmniStorage
- Online access to secondary and tertiary storage through HP's optical and tape libraries
- High-performance data warehouse applications based on HP 9000 Class K and T business servers.

Conclusion

The ideal storage system would provide complete and integrated storage management functionality, smooth integration with multiple file systems and databases across a broad set of operating systems, and support for a large variety of peripherals to satisfy the needs of different storage management applications. Preventing this ideal solution from occurring are a multitude of nonintegrated storage components from different vendors. This situation forces administrators to use single-unit solutions for storage management, which leads to redundant and inconsistent management environments.

The HP storage architecture offers a streamlined and unified interaction among diverse storage components. This architecture allows for customized solutions using plug-and-play components and enables different storage components to interact consistently. For example, HP's backup solutions are being integrated via APIs with databases such as Oracle and Sybase.

References

1. T. Skeie and M. Rusnack, "HP Disk Array: Mass Storage Fault Tolerance for PC Servers," *Hewlett-Packard Journal*, Vol. 46, no. 3, June 1995, pp. 71-81.
2. S. Dimond, "DDS-2 Tape Autoloader: High-Capacity Data Storage in a 5 1/2-Inch Form Factor," *Hewlett-Packard Journal*, Vol. 45, no. 6, June 1994, pp. 12-20.

HP-UX 9.* and 10.0 for HP 9000 Series 700 and 800 computers are X/Open Company UNIX 93 branded products.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

X/Open is a registered trademark and the X device is a trademark of X/Open Company Limited in the UK and other countries.

- ▶ [Go to Article 11](#)
- ▶ [Go to Table of Contents](#)
- ▶ [Go to HP Journal Home Page](#)