

Storage Management: Staying Ahead of the Deluge

By Richard Cole

Better ways have to be found to store increasingly larger amounts of data. It will take the proper technology, planning and cooperation to develop successful storage solutions.

As anyone who works in an office knows, the best way to predict what information you'll need next week is to clean out your desk or your hard disk. As if by magic, the specific files, memos, business cards and phone numbers that you throw away will invariably contain the information that you are looking for.

At the same time, it is necessary to cull dated information on a regular basis and/or retire it to storage. For most organizations, the solution to this familiar dilemma used to be a fairly straightforward matter of manually making tape backups and keeping them in a safe place. But today the task of corporate data storage has been compounded by several factors.

Most obvious is the sheer proliferation of information. Records accumulate, and a history of the organization's activities has to be maintained. New technologies, such as graphics and video, capture more information and consume greater disk space. Furthermore, mergers and acquisitions often force IS departments to manage and store exponentially greater amounts of data as information systems are brought together into single, enterprise-wide environments.

The move from mainframes to distributed client/server environments also has added to the complexity of storage requirements. When almost all data resided on a mainframe supervised in a data center, the focus was on conserving file space and storing data efficiently. The IS manager tried to keep a lid on growth because more storage meant more hardware—an expense that was

charged to the IS budget. But in distributed processing, users have more control over creating and saving information. Because these users are not directly responsible for storage expenses, they are not oriented toward saving space. Often, they simply buy more disks and expect IS to handle storage. The result is booming growth in data storage, often unaccompanied by a consistent policy for storage management.

Many corporations now hold onto their data longer and make it work harder. For example, database marketers are now using demographic information and huge databases to develop an increasingly precise understanding of consumers and buying patterns. This information is sold or brokered to retailers who, more and more, direct their marketing efforts in terms of the "lifetime value" of their customers. These database techniques have helped to expand retailing from mass marketing to niche marketing and now to marketing at the individual level; they also have created a new generation of data management and storage challenges.

Finally, today's users are more demanding about the amount and level of information they want. The development of sophisticated report applications or analysis tools like online analytical processing (OLAP) means that more data can be presented in more ways in a matter of seconds. If the data is not readily accessible or has been stored in an older format that is now incompatible with a corporation's current system, users may protest vigorously. Glenda Lyons, formerly vice president of software technologies at PaineWeber in New York City, tells the story of a stockbroker who could not get the information he wanted fast enough from his desktop computer. His reaction was to throw a steel chair through a 28th-floor window.

The Big Picture

In facing these manifold challenges, systems managers need more than ad hoc approaches. A successful data storage strategy involves several activities, starting with general backup and recovery. This first and most frequent storage activity occurs at the user level on a daily basis. A file, for example, might be downloaded from a PC hard drive to a diskette, then transferred later to a tape storage medium down the hall or at a data center.

Later, this data can be transferred to archive facilities for long-term storage. For purposes of security and climate control, some archives are in former salt

mines located miles underground in Kansas and Louisiana—which gives a literal meaning to the term *data mining*. State and federal regulations require businesses to archive tax, financial and other information for set periods of time. Some regulations are industry-specific. The insurance industry, for example, is required to keep policy information on their clients for the life of the client. In long-term archiving, special attention must be paid to the physical state and degradation of storage media such as tapes, some of which have a “shelf life” of five to 10 years.

Disaster recovery is a related aspect of data storage. This involves storing the information that a business requires to start up again after a calamity. Storage facilities are often kept at safe distances from the main computing site. Files are refreshed more frequently than with true archives, and the storage facility has to be accessible by high-speed communication lines so business-critical data can be brought online quickly if disaster strikes.

How to manage these different short-term and long-term storage requirements prompted the idea of hierarchical storage management (HSM). An HSM solution is a set of migration tools and strategies that enable the continuous retirement and deletion of data based on the frequency of its use. Data is regularly deleted or transferred down the media hierarchy from diskettes and disks to tape and then to cheaper, less accessible media such as optical platters. HSM solutions can be a good way to make data easily available to the user while moderating the storage cost compared with a disk-only solution. In some cases, an HSM solution using short-term disk storage may provide higher access for data that would otherwise be kept on traditional tape systems.

The drawback with HSM is that an integrated, automatic system for storage management that’s suitable for most corporations is still a software generation away. A true HSM solution would have to be integrated with databases, but according to John Camp, research director at the Gartner Group in Stamford, CT, most database management system (DBMS) vendors have not developed the appropriate application programming interfaces (APIs). “There’s no incentive,” Camp says. “They aren’t in the business of data storage management.” He points out that DBMS vendors are more concerned about transactions and keeping the

business running.

In fact, agreement on APIs and software standards in general is needed throughout the data storage industry. The Posix Committee Working Group P1103.1k and the Storage Systems Standards Working Group (IEEE P1244) are trying to develop interoperable storage standards. In addition, the Association for Information and Image Management (AIIM) recently began discussions on the portability of large data archives between heterogeneous data management platforms. During this year’s AIIM convention, users declared that they wanted a way to transfer data from legacy archives without the exorbitant costs of copying gigabytes (GB) and even terabytes (TB) of data to new media. Software vendors, on the other hand, reportedly are still reluctant to agree to standards that might be incompatible with hardware storage technology developed in the future. So far, a broad consensus on interoperability and APIs has not emerged.

Making Do


Lacking a generally available integrated data storage solution, many organizations are developing their own answers; some are fairly specific, others more general and larger in scope. For the narrower focus, network backup is a case in point. The Liggett Group, a manufacturer of tobacco products in Durham, NC, recently implemented an automated tape backup system for its local-area network, using an autoloading tape library. According to Dana Gantt, director of technical services, the company saw an opportunity to move to a fully automated network backup system in 1993 while migrating from an IBM ES/9121 mainframe to a Unix-based client/server environment. Today, the network runs on nine HP 9000 servers with a smaller number of proprietary-based HP 3000 servers, plus 300 PCs connected by a Novell NetWare LAN. Other Unix servers run SCO Unix 3.1.2 (about five years old)

and Sun Solaris. The main DBMS is Oracle.

This migration was prompted mostly by a desire for the efficiencies of using packaged software, but Liggett also wanted to develop an operatorless, “lights out” data center. With the mainframe, several Novell file servers had to be backed up manually using an 8mm tape drive supporting Archivist software from Palindrome of Naperville, IL. The backup was started each night by the second-shift operator. If a tape filled up after the second shift had ended, the backup would not be completed that night and had to be finished by the morning operator. This time lag introduced the possibility of data integrity problems in the backups.

Looking for an automatic backup solution for its client/server LAN, Liggett chose a TLS-4220 tape library from Qualstar of Canoga Park, CA, supporting Palindrome’s Storage Manager version 4.0. The library has two 8mm 8505XL cartridge tape drives from Exabyte of Boulder, CO, and holds 22 tape cartridges for a total capacity of 170GB. Currently, 20 of the tapes are stored in two removable tape magazines, and the other two tapes are stored in fixed slots.

David Channell, systems engineer at Liggett, says that with the new library, he has to load only one set of tapes per week. Once the library door is closed, the subsystem automatically checks its inventory using a bar code scanner. Instead of loading tapes and reading the internal label, the library scans bar codes on the outside of the tape cartridges, which reduces wear and tear on the tapes as well as the tape library and its robotics. With bar-code scanning, the library can also learn more quickly which tapes are currently loaded in the tape library. Channell and his colleagues do not have to maintain external labels, and in the event that staff members need to read a label, they can do so via the Palindrome software.



The tape library uses a Tower of Hanoi tape rotation scheme, based on a sequence of moves from a popular mathematical puzzle. The Tower of Hanoi pattern uses less media than other rotation patterns and allows files to be changed with a variety of file versions. "Essentially, you can store more stuff on fewer tapes," says Channell.

Out on the Edge

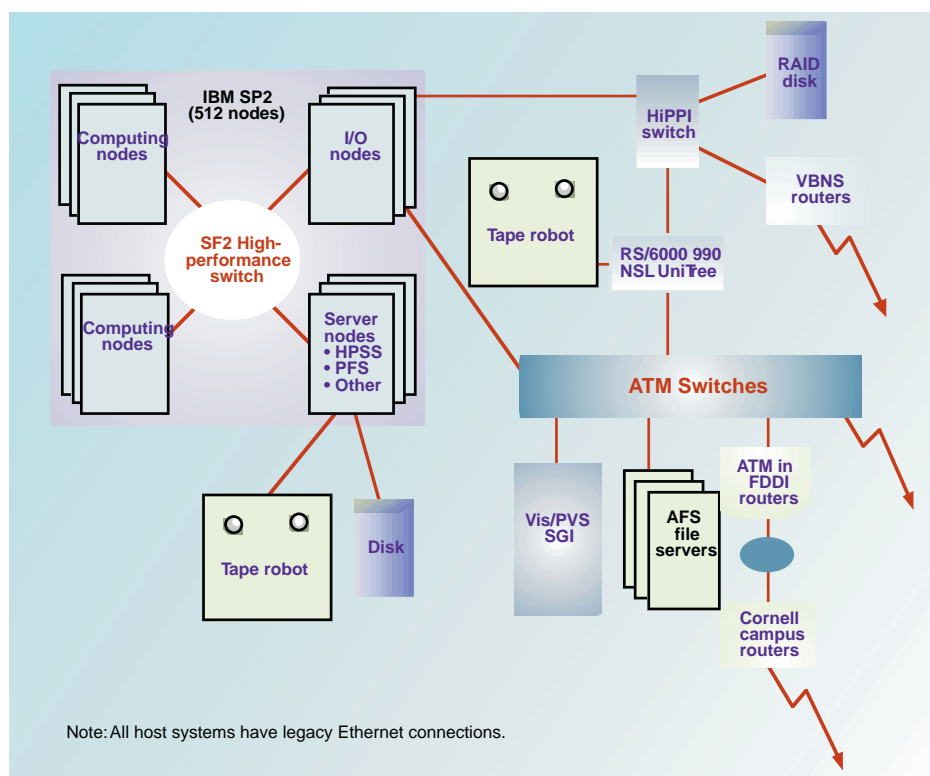
The Liggett tape library regularly backs up about 20GB of data, a typical amount for many corporate environments. In contrast, massive systems dealing in terabytes of data have been implemented at several academic supercomputing centers and national laboratories across the United States. These large systems not only represent the latest in technology, they offer a glimpse of the future for many corporations as data storage systems continue to grow.

The Cornell Theory Center at Cornell University in Ithaca, NY, is the sixth largest computing center in the world. It provides supercomputing services to academic and corporate researchers across the country. Data comes from a variety of sources ranging from the radio telescope at Arecibo, Puerto Rico, to electron microscopes.

At the heart of the center is an IBM SP2 supercomputer supporting 512 processors in a massively parallel processing (MPP) environment running AIX, IBM's Unix variant. The center also runs several supercomputers and high-end workstations for visualization, including a Power Visualization System (PVS) and three Onyx computers from Silicon Graphics, Inc. (SGI), and supports about 200 Unix workstations, 50 Macintoshes and a few other PCs.

Connectivity is especially important to operations at the center, since 60 percent of users are located at other facilities. Typically, researchers access data and run programs in terminal sessions from their own sites. The Internet is the main method of communication between them and the center. The center is connected with NYnet, its New York asynchronous transfer mode (ATM) network, and NYsernet, an Internet service provider. In addition, there is an experimental network called VBNS (very high speed backbone network server), which provides a dedicated ATM line running at 155Mbps between Cornell and other supercomputer centers.

Unlike businesses, the center is not legally required to hold data. Files are automatically backed up and stored during research projects, but they are overwritten once a project is completed. Users



The Cornell Theory Center, as of 1995, featured supercomputers, high-end servers and various workstations whose storage needs were addressed through high-speed switches and robotics.

have ultimate responsibility for their own backups. However, the massive amounts of information handled by projects at any one time require an equally massive storage solution. Using the Andrew File System for parallel file serving, the center transfers data to two IBM 3494 tape robots, each capable of holding 1,500 tapes and 15TB of uncompressed data. These libraries are accompanied by 10 IBM 3590 Magstar tape drives with a capacity of 10GB per tape.

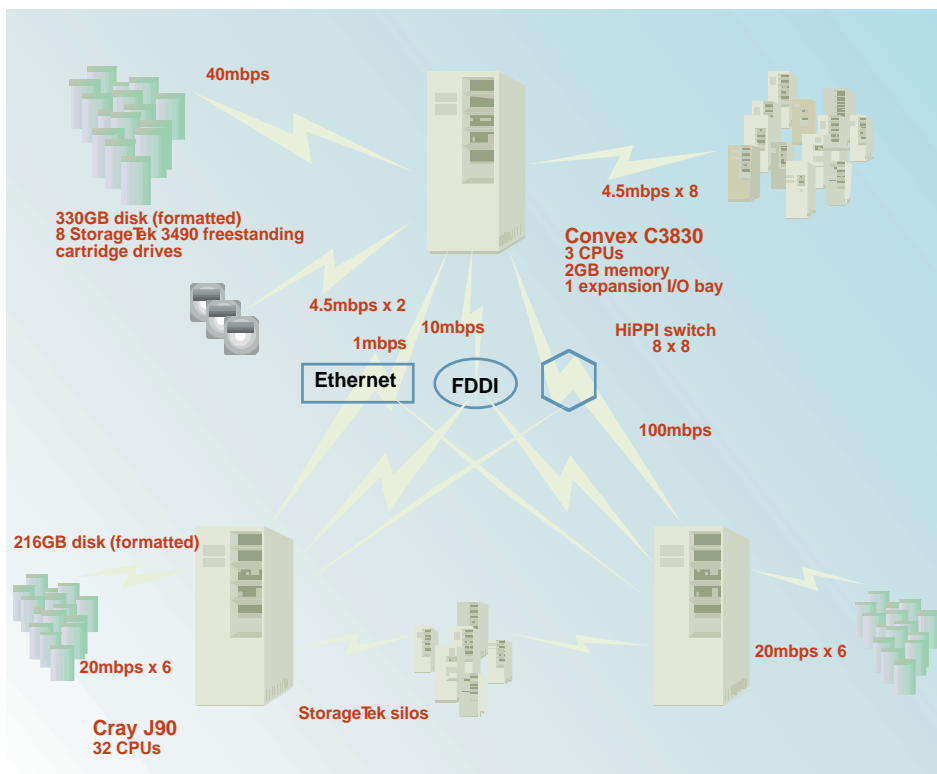
Doug Carlson, associate director for systems and operations, points out that the center has several unique features for handling mass storage. One is the High Performance Storage System (HPSS). This technology, according to Carlson, represents one of the highest-performing mass storage systems available today. Developed by IBM's government systems division and four national laboratories, HPSS is designed to provide a highly scalable parallel storage system for MPP systems. In this context, scalability includes data transfer rate, storage size, number of name objects, size of objects and geographical distribution. "When fully implemented this year, it will allow us to get data rates much higher than before," says Carlson. HPSS has been built to hold billions of directories, billions of files and petabytes of data. (A petabyte is a quadrillion, or

1,000 trillion, bytes.) The HPSS works with a parallel file system for the SP2.

In the future, the center will continue to expand its processing and storage capacity, but to do so it must overcome more technical challenges. In the industry generally, Carlson says, processing capacity is growing 50 percent per year, while I/O capacity is only increasing at 20 percent per year, creating a bottleneck. As a solution, he suggests several possibilities, including parallel file systems and parallel tape support. He also mentions implementing faster networking systems based on ATM and High-Performance Parallel Interface (HiPPI) switching.

Massive Requirements

The mass storage and scientific computing branch of the NASA Center for Computational Sciences (NCCS) at Greenbelt, MD, provides another example of massive storage on a grand scale. Headed by Nancy Palm, the NCCS supports research and modeling for the earth and space science community. The NCCS runs two Cray J90 supercomputers and one Convex server. Primary access for users is through Ethernet, with Fiber Distributed Data Interface (FDDI) and local HiPPI support between the supercomputers. Users run a variety of Unix-based IBM, Sun, SGI and other workstations.



The NCCS network combines supercomputers, high-speed connections and various storage systems.

As would be expected, the NCCS shares some of the technical solutions found at Cornell, in particular an HSM solution based on UniTree. Developed at Lawrence Livermore National Laboratory in Livermore, CA, and owned by UniTree Software of Dublin, CA, UniTree provides service software and coordination to transfer very large amounts of data. For data-intensive environments, one of its more important features is its "virtual disk" technology. Files are available to clients in the form of a disk with apparently unlimited capacity. The UniTree manager automatically migrates infrequently accessed files away from the high-speed disk cache toward the tape media. When an archived file is requested, the software manager automatically restores the file. The user does not need to know the physical location, and storage space is almost unlimited.

Palm explains that about 194GB of data are held online on the Convex server dedicated to the UniTree disk cache. Another 34.24TB are robotically managed "near-line" on the Convex and Crays. The next level comprises seven 4.8TB "silos" and one 1TB Wolfcreek silo from StorageTek of Louisville, CO. These silos are massive robotic tape storage units; six are managed by Convex UniTree for permanent storage, and two are used for Cray short-term storage. There is also an operator-mounted,

offline tape archive unit with 4.3TB of vaulted storage managed by UniTree.

Massive data storage is especially important to her environment, Palm says, because many of her users constantly access the same data and add more data to it to develop, say, a model of weather patterns. Data 50 years old is just as valuable as data gathered last week. In developing an ever-enlarging, long-term storage system for these scientists, she stresses the importance of maintaining the future readability of data through standard storage media formats as well as metadata that provides the names and locations of files. She also points out that database architectures must scale as much as possible. A growing parallel environment may require an equally parallel database architecture. These challenges may seem to be beyond the scope of today's commercial environments, but perhaps they are not for tomorrow's.

Beyond Technology

For general advice about implementing or expanding a data storage system today, the comments of IS professionals are remarkably consistent. First of all, everyone stresses the importance of long-term planning that includes data storage. "Many traditional capacity plans have focused mainly on CPU capacity, yet data storage capacity is just as important," says Carl-

son of Cornell.

Don Crouse, vice president of technology at Large Storage Configurations, a vendor based in Minneapolis, says that data storage planning should be as specific as possible. "You have to ask yourself how much money you want to put into storage management and then create a plan that answers this question in terms of storage space, retentivity [how long the data is held] and the resultant cost."

Secondly, everyone agrees that proper storage management is an intensively collaborative, cooperative process. IS has to plan internally, but this plan has to be based on talks with vendors, management and users. "We look for a cradle-to-grave relationship with vendors," says Palm. "Customers and vendors need to be honest in discussing future requirements. And not in a finger-pointing way, but working 100 percent together."

The relation with management is equally important. Joshua Greenbaum, senior analyst for Sentry Market Research in Westboro, MA, cites as an example a large, international car rental agency that wanted to build a central database of information from all its affiliates in the United States and overseas. The project almost crashed because the affiliates had not fully agreed to back it. "Decentralized data means decentralized data management," he says. This makes it necessary that management be properly informed, educated and sold on data management and storage projects. "Technology is important, but a lot depends on cooperation, and that has to be done by good old-fashioned interpersonal relationships."

In a similar way, talking with users is a critical factors for success. Palm says that every five years, the computing environments and research requirements committee asks scientists working with NASA about their requirements for the near, middle and long terms. "These guys know what they need, so as part of a service organization, we take their requirements and turn them into [procurement] documents to make sure we meet them with the right hardware, software and manpower."

Based on current trends, data storage will only become more demanding. To succeed, IS departments will have to have a good idea of where they need to go, based on a continuing dialog with the people who develop, support and use the systems. **IT**

Richard Cole is a contributing editor to *UniForum* publications. He can be reached at 76402.1503@compuserve.com.