RHEINNER GUIDE

Automated Tape Storage for Document Imaging

Document Management, Imaging, and Workflow Technology Guide Series

This guide has been sponsored by

StorageTek

StorageTek… where the world's information goes.

StorageTek Technology Corp. designs, manufactures, markets and services, worldwide, information storage and retrieval subsystems for enterprisewide computer systems and networks.

StorageTek are the creators of the Nearline™ architecture, an innovative combination of automated tape libraries, high speed and high capacity tape drives, enabling software and intelligent microcode. The Nearline architecture enables random access to long-term, high-capacity, low cost, document and image archives for the department or the enterprise.

Solutions can be implemented at either the department or enterprise level. They can be shared among applications and host platforms. The open architecture of Nearline enables access to documents and images throughout the organization without compromising security.

Alliances have been formed with industry leaders to provide world-class document management solutions including:

- micofiche replacement
- online viewing
- report management
- imaging
- workflow
- process re-engineering

StorageTek has offices worldwide including over 80 offices across the U.S. Contact your local account team for more information, or call 1-800-234-5290.

StorageTek

Visit our web site at:
http//www.storagetek.com

Table of Contents

The Rheinner Group

The Rheinner Group is a leading research, consulting and education firm in the document imaging, management and workflow industry. Its Certified Document Imaging Architect (CDIA) Education Program, which covers many of the same issues addressed by The Rheinner Group's Technology Guides, is the most popular training program in the imaging industry.

For more information on The Rheinner Group, CDIA course schedules, or to obtain help designing and implementing document imaging and workflow systems, please call 617-741-8100. Also, please visit our web site, The Rheinner Café, at www.rheinner.com.

The Rheinner Group is a proud sponsor of the Certified Document Imaging Architect (CDIA) certification program.

Introduction

Document imaging applications place unique demands on the storage subsystem. First, document images have very different life cycles than do character-based data files. While character-based files typically have a short, but intensely-active life, document images must often be kept in digital form for an extended period of time. In addition, since document imaging is aimed at making paper-intensive operations more efficient, there is generally a substantial volume of document image data. The end result can be millions of image files, some portion (or all) of which must be accessible at all times.

To accommodate these characteristics, a document imaging storage subsystem must deliver high capacity, data reliability over extended time periods, and rapid access to data. Because no single storage technology can simultaneously meet all of these performance requirements, the storage subsystem must be configured with a variety of storage devices. For example, while the speed of fixed magnetic disks provides rapid access to data, their limited storage capacity is not well-suited to long-term document imaging storage and archival needs. And, while robotic storage devices offer high capacity, media removability and very low storage costs, the technology does not offer the access speeds necessary for many imaging applications. Yet, when fixed magnetic disk is combined with a robotic storage device, the subsystem can deliver the capacity, reliability, and access performance required by document imaging.

Other types of storage devices, including tape drives, WORM (Write Once Read Many), or CD-ROM, are frequently included in the document imaging storage subsystem, as well. For many environments, these devices can provide data storage, backup, or access as cost-effectively as traditional optical storage technology. One example is the automated tape library, which can deliver the same, and in some cases more, functionality as optical disk jukeboxes for storing and retrieving large volumes of data, at a much lower cost.

The virtually inexhaustible capacity of tape libraries, coupled with the unquestionable economies of tape storage over any other storage media, make it a desirable choice for applications that generate as much data as document imaging solutions. Although commonly overlooked, tape represents an important functional element of the document image storage subsystem. Traditionally considered a backup or far-line technology, automated tape libraries have not enjoyed the same attention paid to the fixed magnetic disk industry. And yet, the progress made by the tape industry has been substantial. Sophisticated tape storage libraries are able to store vast amounts of data on internal shelves. And, the flexibility offered by automated tape libraries can offset some of the design issues associated with the unpredictable nature of document imaging media requirements. With continued advances in tape drive performance and its undeniable economic advantage, automated tape storage can provide a configuration flexibility unrivaled by other media.

This Rheinner Group Technology Guide describes tape storage and automated tape libraries. After outlining the features and functionality provided by automated tape libraries, the guide explores the role of automated tape libraries in a document imaging environment.

What is Tape Storage?

At its simplest, tape storage refers to the recording, storage, and retrieval of data on magnetic tape. Magnetic tape comes in a large assortment of lengths, widths, and capacities. Small tape cassette packages can provide approximately 2 gigabytes (GB) of storage capacity while 1/2″ and 3/4″ videotape-style media can hold up to 50 GB of data per tape.

IBM first popularized the notion of tape storage with the introduction of systems in 1953. Since then, tape storage has continually demonstrated its compelling economic advantage as a high-capacity, long-term medium for data storage. On a pure media cost basis, tape is the most cost-effective storage medium available. In other words, it is cheaper to store data on tape than on any other medium, including paper, or fiche. For example, the cost to store one million document images on tape is about $70. With the $3 price of 500 sheets of 8.5 x 11″ paper and office space rates over $15 a square foot, storing information on tape is an order of magnitude cheaper than storing the original paper document. And, tape has proven its ability to comfortably house data for over 30 years, with no significant data loss.

Tape Form Factors

Tape media comes in many varieties and goes by many names. The most common measurement and designation refers to the width of the tape, for example 4mm or 1/2″ tape. In addition to traditional reel-to-reel tape, other common tape form factors include:

- QIC Data Cartridge: 1/4″ tape, unique packaging system;
- IBM 3480/3490: 1/2″ tape in a "standard" package;
- DAT: digital audio tape, using 4mm digital audio and 8mm video tape; and
- D-1 through D-6: video tape, using either 1/2″ or 3/4″ tape that can also be used to record data.

A number of vendors manufacture transport mechanisms or drives to read and write data to and from the media listed above. While there are many ways to record data on tape, the two primary methods are digital linear tape, and helical scan recording. Linear tape devices, which use the same format as the IBM 3480/90 cartridge, can deliver between 200 megabytes (MB) and 20 GB of data storage capacity. Helical scan devices, although more expensive than linear tape mechanisms, can record up to 50 GB of data on its cartridge media.

Tape Performance

Because tape is a linear medium, data is located along a continuum and can only be located by searching either forward or backwards. There are a number of ways to bring efficiency to this process, the most common of which is to memorize data locations on the tape and immediately forward to that portion of the tape. Capacity is added by recording on more than one track (such as 18 and 36 track recording), and by using longer tape. Because tape is linear and not a disk-shaped media where any particular data location is easily found from the disk center, the access speed (the speed at which it finds the data location) is always slower than for other media.

However, tape's data transfer rate is impressive: A helical scan tape drive can transfer data in excess of 11 MB per second, nearly double the data transfer rate of most hard disk drives and ten times as fast as optical media. Additionally, the ability to store vast amounts of data very inexpensively on tape, coupled with tape's reliability, make it an attractive storage option for applications with high-volume storage requirements.

The Storage Hierarchy

Storage components and media can be classified according to a hierarchy based on the speed, cost, and capacity offered by each technology (see Figure 1). By utilizing the variety of devices and media according to their position in this hierarchy, users with intensive storage requirements, such as document imaging users, can achieve the optimal balance of performance and cost.

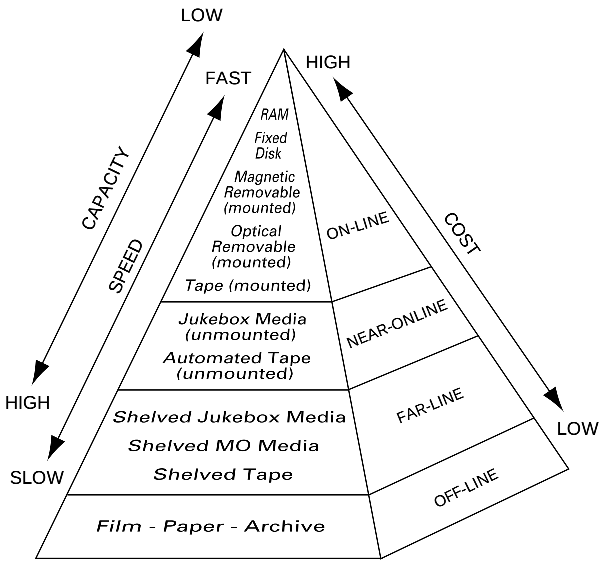

Figure 1: The Storage Hierarchy

Computer storage of data is generally accomplished using a combination of three primary types of media.

1. Solid State

Random access memory (RAM) chips are inserted into the processor board of the computer. The fastest data storage media in terms of response, RAM is also the most volatile. When power is interrupted, all data is lost. Because of its volatility and finite storage capacity, solid state's speed is most often reserved for the computer's microprocessor, storing only the data immediately required by the microprocessor.

2. Magnetic Storage

Magnetic media is characterized by its ability to reliably and quickly store large amounts of data for short periods of time. There are two principal magnetic storage devices:

- Fixed Disk Drives—hard disk drives in which the magnetic media and the drive itself form a single unit. The storage media cannot be removed without also removing the drive.
- Removable Media Drives—drive mechanisms in which the magnetic storage media (either disk or tape) can be removed without removing the drive mechanism.

Removable disk media consists of a magnetic disk media that is housed in a protective plastic casing. Available in a range of sizes ranging from 3.5-inch to 5.25-inch media, capacities for removable disks range from 1.4 MB for floppy disks to 1 GB for a JAZ cartridge. The variety of media and packaging options make most high-capacity removable disk media non-interchangeable with other devices even in the same size form factor. Thus, SyQuest 3.5-inch cartridges can only be used in SyQuest 3.5-inch disk drives, and 3.5-inch ZIP disks can only be used in 3.5-inch ZIP drives.

Removable tape drives utilize removable magnetic tape, rather than disks, to store data. The interchangeability of tape media is also highly dependent on form factor and drive manufacturer. For example, even though the media size is the same, not all 8mm tapes recorded on one device can be read from and written to any other 8mm drive. The industry standard 3480/90 format is an exception, and can be read by any 3480/90 device.

3. Optical Storage

There are two principal categories of optical storage: WORM and MO. WORM optical storage devices can read from WORM media numerous times, but can only write to the media once. Data recorded to WORM media cannot be overwritten; instead, data is recorded to a new portion of the disk. While WORM media can be used repeatedly to access previously-recorded data, the ability to record new data on a particular WORM disk is finite. WORM storage capacities ranges from 600 MB to 14 GB per disk or platter, depending upon the size of the disk and the storage device (drive or jukebox).

CD technology is related to WORM technology, but performs recording functions in a dramatically different fashion that impacts how CD can be utilized.

The performance of each media type is based on the inherent characteristics of the media itself, as well as the features of the device used to record and retrieve data from the media. While the functionality of each device and its associated media performance is a critical issue, another key dimension of the storage hierarchy involves the logical location of computer data. In order to select the appropriate devices for the imaging storage environment, the document imaging architect must determine the movement of documents throughout the processing cycle. The storage subsystem should be designed to optimize data availability so that data is stored economically, but is available within a reasonable timeframe. All storage systems, indeed the imaging system architecture in total, are built around the logical location of the data. The logical location of data can generally be classified into one of four categories.

On-Line

Data that must be immediately available is typically stored on media and devices capable of immediately responding to the request. Internal magnetic media constitutes on-line storage, as does data on a removable disk provided that the disk containing the required data is mounted in the device.

Near-Online

Data that requires less-frequent access is typically stored near-online, where some intermediate media manipulation is required before the request can be serviced. Currently, the most common robotic device in document imaging is an optical disk jukebox that contains storage shelves for disks. A robotic arm retrieves the appropriate disk and inserts it into internally-mounted optical disk drives so that the data request can be fulfilled. Data in a jukebox can be on-line if the requested data is on a disk that is inserted and mounted in one of the internal drives. For this reason, a tape or optical jukebox is also commonly referred to as providing near-on-line data. Data located on another disk or tape housed in the jukebox is considered near-online.

Far-Line

Data is considered far-line when it is still available to the computing device but requires considerable time for location and retrieval, such as tape or optical media that has been placed on a shelf outside the jukebox.

Off-Line

Data is considered off-line when data has been completely removed from the storage system, and requires a great deal of manual intervention (usually a human operator) to locate the archive (which may be in a remote location), retrieve the data, and make it available to the system. Off-line storage is typically reserved for data that must still be kept, but does not need to be retrieved often or immediately.

Document images which must be located quickly need to be kept on devices listed near the top of the hierarchy pyramid. Document images that must be available for occasional access should be stored on near- or far-line devices, or in other words, accessibility should be determined by the speed at which the data is required to be available in the event of access. Images that are rarely or never retrieved should be migrated toward the bottom of the storage hierarchy, onto off-line storage devices, or deep archival storage media.

Once the type of storage technology is determined, a series of specific storage products must be selected. This selection is based on performance, capacity, reliability, and price; the importance of each will depend on the primary function the device will be called on to perform.

The most problematic storage decisions typically occur when considering storage alternatives for near-online data availability. When large volumes of data must be retained for extended time periods but must also be available on request, the architect must deal with the tradeoff between speed, capacity, reliability, and price. No device can provide the speed of hard disks and long-term, economical storage of large volumes of data. Instead, users look to optical disk jukeboxes and automated tape libraries to adequately and cost-effectively meet these storage requirements. While optical jukeboxes and tape libraries provide similar storage services, there are some notable differences.

The Automated Tape Library

Like optical disk jukeboxes, automated tape libraries are robotic devices. A robotic mechanism retrieves, delivers and inserts a tape cartridge into an internally-mounted tape drive. The automated tape library can have up to several of these drives (which are also referred to as tape transports). In computing vernacular, every device capable of reading media is called a “drive”. Some tape library manufacturers use the term drive when referring to the entire device, including the recording/playback device, cabling, and power supply. These manufacturers refer to the individual recording/playback device as a transport. In this case, a single drive unit can have multiple tape transports. In this guide, the term tape drive and transport are used interchangeably, and both refer to the actual mechanism into which the tape is inserted.

Automated tape libraries come in a variety of shapes and sizes, with an assortment of available features. Automated tape library models can accommodate from 20 to 6,000 tape cartridges, which are most often housed within the library unit itself. Some house just one media type; others can contain several different types of media, with varying capacities. There are also a number of different methods to configure the library storage module (LSM), which contains the tape cartridges. Certain models allow LSMs to be interconnected via conveyor belts to provide additional storage capacity. The cartridges are moved from one LSM to another on these conveyor belts. It is also possible to have a LSM without any internal tape drives at all.

When configured within an automated library, tape storage’s favorable economics become even more compelling. With multiple tape cartridges in a single library housing, the cost of the library is amortized over a large volume of data. The larger the body of data, the better the economics. The economic benefits are even more significant if the automated tape library is used to both store and retrieve data. Consider archive retrieval on tape vs. film. If data is stored on film, a human must be involved in every retrieval, which is considerably more expensive than if data can be automatically retrieved from a workstation. For this reason, the automated tape library is sometimes the preferred medium for computer output to microfilm (COM) or COLD (computer output to laser disk) applications.

The capacity of an automated tape library depends on the number of cartridge slots and the type of media utilized. At the low-end are libraries that hold 20 tape cartridges and deliver capacities of about 60 GB. High-end products can be classified into libraries that hold between 200 and 500, 500 to 1,000, and 1,000 to 6,000 cartridges. These high-capacity, high performance devices can be configured to deliver hundreds of terabytes (TB) of data storage capacity. This potential capacity is unlikely to be exhausted in a lifetime, making the automated tape library a cost-effective choice for applications with very large data storage requirements, such as document imaging.

Management

Devices capable of supplying so much data on a near- and on-line basis require extensive data management and file organization. As with any removable media device, automated libraries must be managed for effective and optimal performance. The major management issues are:

- managing data location by slot, and cartridge;
- prioritizing requests for exchanges in order to avoid unnecessary trips for the robotics;
- managing an available pool of tape media;
- managing device allocation to meet data access and performance requirements;
- managing tape reinventory in case of moved or switched cartridges; and
- refreshing tape media, such as retensioning tapes which have been idle for a long time and rerecording old data onto new media.

These functions are available in software provided by the automated tape storage manufacturer and from third parties.

Performance Criteria

Robotics located inside the LSM perform the retrieval and mounting of the tape cartridge in a tape drive. Most robots can only carry one cartridge at a time, which means that only one robotic operation can be performed at a time. An optical jukebox works this way: the robot makes one trip to retrieve a previously-read platter from a drive, puts that platter back into its storage slot, takes another platter out of its slot, and lastly, inserts the newly requested platter into the drive. Automated tape libraries, on the other hand, dismount tapes as soon as the request has been fulfilled, provided there are no other requests pending for that tape. This eliminates the dismount procedure required by optical jukeboxes before a mount can be fulfilled. Automated tape libraries retrieve a requested tape from its storage slot and place the tape immediately into a drive. Some tape library manufacturers have developed robotic devices that can carry two cartridges simultaneously. These robots skip part of the previous sequence by picking up a second cartridge at the same time as the first, or by dismounting a cartridge out of each of two tape drives on the way back to the storage slots.

Automated tape library robotics are typically rated according to the number of exchanges per hour (EPH), which is the number of times the robot can remove a tape from the tape drive, return the tape to the shelf, retrieve another tape cartridge, and insert it into the drive. Robot EPHs vary considerably; the fastest robots are able to perform more than 350 exchanges each hour (or one exchange every ten seconds), making them competitive with any robotic storage device.

Obviously, multiple drives can improve the responsiveness of the automated library considerably. The ability to keep several volumes of data mounted reduces the need to perform exchanges, which results in dramatic savings in retrieval times. The number of drives that can be kept in an automated library ranges from one to 16. Naturally, if the storage library itself must handle mixed media, then more than one kind of drive must be available.

Access speed, while an important performance measurement for fixed magnetic and optical disk drives, is somewhat irrelevant for magnetic tape devices. Access speed is the amount of time, measured in milliseconds, that it takes for the magnetic or optical reading device to position itself directly over a portion of data that has been requested. Because tape is linear and not disk-shaped, the time it takes to find a particular bit of data depends on the length of the tape, the recording density used, and the particular position of that data. In other words, the access speed varies, making the measurement somewhat meaningless.

Instead of access speeds, tape drive manufacturers differentiate themselves on the data transfer rates of their devices. The data transfer rate measures the amount of data that can be transferred to or from the recording media on either a sustained or burst basis. Dependent on the drive mechanism, tape library data transfer rates range from 3 MB per second to 11 MB per second. The low-end of the performance spectrum (3 MB per second) is on par with hard disks while the highest-performing drives (11 MB per second) deliver twice the data transfer rate of most hard disk drives. And, tape drives boast a higher data transfer rate than most optical disk drives, which generally top out at 3 MB per second.

A high data transfer rate is critical for transferring large amounts of data from near- to on-line media, for example, when document images that need to be processed are moved from tape to magnetic disk. Similarly, when these images are moved in bulk back to magnetic tape, its high data transfer rate can enhance the performance of the entire system.

When considering the performance of tape, one final attribute to consider is its long-term reliability. High-quality tape cartridges offer data storage reliability of about 30 years. Additionally, given the low price of tape media, data can be inexpensively backed up by configuring the automated tape library to rerecord data on a continual basis.

Performance Dimensions

With a multitude of capacity, drive and robotic options available, automated tape libraries can be configured to meet a broad range of performance requirements. Optimized tape library performance becomes a function of specified library and media capacity, robotic exchanges per hour, drive mechanism access, and data transfer rates.

Capacity

Libraries can be selected on the basis of required capacity, with anywhere from 20 GB to 300 TB of data storage capacity.

Drives

A range of tape drive configurations are available, including helical scan technology tape drives that can record up to 50 GB to high-capacity linear tape drives capable of storing from 200 MB to 20 GB of data. Drives can also be selected on the basis of their access speeds and data transfer rates. Finally, the need for mixed tape media in order to maintain backwards capability can impact drive selection. It is important to point out that tape formats have been backwards compatible since they first came into use in the late 1950's.

Robotics

Tape library robotic speeds range from 90 to over 350 exchanges per hour. Depending on the amount of media exchange anticipated for the application, users can select libraries that deliver the appropriate EPH figures.

While optical disk jukeboxes can also be configured to deliver a range of capacity and functionality, optical disk drives do not offer the same range of performance. Nor, do optical disk jukeboxes offer the same range of options for the robotic mechanism's performance. Depending on its configuration or capacity, an optical disk jukebox can take between 15 and 25 seconds to deliver data that resides in the jukebox. An automated tape library could perform the same activity in anywhere from 15 to 90 seconds, depending on its configuration. A great many factors related to the actual implementation, the software, the application, and the workload could sway performance either way. In many cases, the automated tape library provides comparable performance to the optical disk jukebox, but in other ways could provide even better performance by also providing a vast storage repository as well.

Subsystem Configuration Options

The net effective performance of an automated tape library is determined by:

- the robotic EPH;
- the number of tape cartridge slots;
- drive performance; and
- media capacity.

Libraries can be optimized for specific application environments and cost/performance tradeoffs can be made. For example, in a large archival library where vast amounts of near-online data storage is required, a robotic device might be paired with a large number of slots and media with substantial data capacity. In this case, an exchange cycle of 60 seconds from data request to on-screen delivery may be less important than having terabytes of near-online storage. Or, the library could be configured to provide terabytes of storage, as well as below 20 second exchange cycles if a device with faster drives and robotics is selected. This configuration flexibility is one of the greatest advantages of automated tape libraries. Depending on the performance requirements of the application, any number of parameters, at both the device and the management software level, can be adjusted to yield positive results.

Further performance flexibility is provided by pairing an automated tape library with other types of storage devices. This union is common in document imaging applications, where the optimal storage environment balances storage economics, data availability, and system response. Since tape, optical, or fixed magnetic devices cannot alone deliver all the features required by a document imaging application, a storage subsystem that blends the best of each is typically implemented. For example, magnetic hard disks are often used to buffer the movement of the robotics found in high-capacity devices, such as optical disk jukeboxes and automated tape libraries. New data coming into the system can be placed on magnetic hard disk before it is filed onto the long-term storage media. Additionally, stored data that is required for processing can be pre-fetched from the repository and placed onto fast magnetic storage so that it is readily available to users. This solution optimizes system performance and storage economics.

Storage Requirements in Document Imaging Environments

The expanding volume of image data associated with many document imaging applications can quickly turn into an expensive and formidable storage challenge. To balance performance with storage economics, the flow of image data must be classified according to its logical location. Once the volume of image data that must be located on-line, near-online, far-line, and off-line is identified (or estimated), storage peripherals that accommodate that availability should be identified and evaluated. A critical piece of this evaluation is to analyze system capture and retrieval patterns, as these activities place the greatest demands on the storage subsystem.

As a general rule of thumb, the file size for a compressed document image is approximately 50 kilobytes (KB). While each document produces a different size compressed image file that results from the amount and type of information on that document, storage requirements eventually average out to approximately 50 KB per image. Depending on the application, between 1,000 and millions (or more) images may arrive for daily scanning into the document imaging system. That means that scanning can add anywhere from 100 MB to 5 GB of new image data to the system each day.

For high-volume applications in which the daily scanning load approaches 5 GB or required storage capacity, the storage subsystem must be carefully configured and managed to accommodate daily capture and processing requirements. In some cases, image data must be accessible 24 hours per day. In such an environment, the storage system must support a constant flow of data from fast, but expensive, hard disks to reliable, economical storage media that still can be accessed directly.

Foldering vs. Transaction Processing

The design of the storage subsystem will be driven by the type of document imaging application, as well. Most imaging applications can be classified into one of two categories. The first are foldering applications built around the file-folder paradigm. The file-folder paradigm maintains all documents related to a particular case, person, or activity in one single, or a series of interrelated file folders. Examples of folder-based applications are litigation support, in which all of the documents pertaining to a particular case are grouped together; or medical records, in which all of the documentation about a particular patient is kept together. Typically, access is relatively unpredictable — any number of documents and folders could be required at any time. Access is often driven by external events, such as the filing of a suit or the admission of a patient.

Transaction-processing applications, on the other hand, generally are characterized by more predictable image access, capture, and retrieval patterns. Examples include the processing of new credit card applications and processing monthly bill payments. Designed to quickly process a large number of transactions, transaction-processing applications may or may not require access to the document image once the transaction is complete. The image can either be kept in a digital archive, or on film. In many cases, access is frequent enough to make microfilm an uneconomical alternative. In this case, maintaining the image in digital form can significantly reduce the cost of retrieval. The use of an automated tape library for these types of applications can greatly reduce the cost of retrieval and storage, and greatly improve response times.

Storage Management Techniques

Regardless of the application type or the volume of image data handled, the vast majority of document imaging applications benefit from the use of one or more of the following storage management procedures.

Caching

Caching is the temporary storage of active documents until they are forwarded for processing or moved to their permanent storage location. In imaging environments, cache is typically fixed magnetic storage. Document imaging storage solutions make widespread use of caching to buffer robotic retrieval, and/or to handle large amounts of incoming data from scanning operations. This is accomplished by placing magnetic repositories throughout the storage subsystem wherever the active document traffic is at its heaviest.

Pre-fetching

Pre-fetching is the predicted movement of document images. It is primarily utilized to move data required for a processing interval from archive, or from other slower access storage facilities, to the magnetic cache set up for active document processing.

Migration

Migration is the automated movement of files throughout the storage infrastructure. Document imaging systems make extensive use of file migration in order to move documents to and from active cache, to pre-fetch files, and to move rarely used files to archive.

Automated Tape Storage in Document Imaging

Automated tape libraries can provide numerous benefits to a document imaging solution, providing near-online storage support, data backup, and/or film replacement. One major attribute is the automated tape library's range of configuration options. A tape library can be configured to hold a month's work on a single tape. Data can be directly accessed in under 30 seconds or transferred to magnetic cache overnight. Alternatively, automated tape can provide instantaneous backup to any of the other storage devices utilized in the system. It can even be optimally configured as both of these alternatives.

For most document imaging systems, the main data repository is a near-online device that can house large volumes of data, and deliver any portion of that data to the user in under 30 seconds. While this is typically associated with an optical disk jukebox, tape libraries provide the same function. In some ways, tape libraries can exceed jukebox functionality by providing more storage capacity, and by utilizing the faster data transfer rates found in certain tape drives at a much lower cost per megabyte.

When using 2.4 GB media, for example, the typical exchange, mount and data delivery cycle is on the order of 23 seconds. Between 6 and 10 seconds of that time is allocated to robotic movement; the remainder is the amount of time it takes to locate the data. This response is comparable, give or take a couple of seconds, to the performance of an optical jukebox, depending on the organization of the data. Additional tape drives could be using 50 GB media, and have that data available for access without requiring a robotic mount and dismount. Additionally, the high data transfer rate of tape drives means that a large volume of data can quickly be transferred from tape to magnetic cache. Because most optical disk jukeboxes can only hold 1.3 GB of data in a drive at any one time and because of their slower data transfer rates, it will take much longer for a jukebox to move the same amount of data to the magnetic cache. In this case, the automated tape library represents a superior solution to the jukebox.

The use of large amounts of RAID (Redundant Array of Independent Disk drives) storage is becoming popular for many imaging applications. The RAID devices typically house all of the active documents, while also acting as a buffer to robotic near-online devices. The use of RAID is gradually lowering the direct access requirements for near-online devices, making the ability to transfer data quickly from near- to on-line the more important benchmark. Under this scenario, the near-online capacity and data transfer rates of tape drives are becoming more important design considerations than the ability to directly deliver on-line data.

Backing up document imaging data is problematic due to several factors. First, imaging systems tend to be very busy during processing cycles, often leaving no time when the system is free to perform a backup. Second, the volume of data may not permit a backup to be performed within a 24 hour day. For example, a system taking in 25,000 images a day may require more than 24 hours to make a copy of that much data. For busy document imaging systems, the amount of time required to perform backups may simply not be available. Therefore, the manner in which the backup will be performed must be an intrinsic part of the storage system architecture. Given the high data transfer rate of certain tape transports (in excess of 6 MB/sec.), and the fact that most libraries can house several transports, data backup can be automatically performed as part of the library's function.

The ability to add very inexpensive media ($7 for 2.4 GB, $70 for 50 GB) on a continual basis is another strength of the automated tape library. The fact that the media is reliable enough to store data for 30 years makes it a highly desirable storage option on a capacity and reliability basis alone. For this reason, automated tape libraries are justifiable even as a replacement for film. Transaction processing systems that require the use of document images for the processing work cycle and then save to film, can realize significant savings by using an automated tape library for active processing and archival, rather than converting to film.

The virtually inexhaustible capacity of tape libraries, coupled with the unquestionable economies of tape storage over any other storage media, make it a desirable choice for applications that generate as much data as document imaging solutions. Although commonly overlooked, tape represents an important functional element of the document image storage subsystem. Traditionally considered a backup or far-line technology, automated tape libraries have not enjoyed the same attention paid to the fixed magnetic disk industry. And yet, the progress made by the tape industry has been substantial. Sophisticated tape storage libraries are able to store vast amounts of data on internal shelves. And, the flexibility offered by automated tape libraries can offset some of the design issues associated with the unpredictable nature of document imaging media requirements. With continued advances in tape drive performance and its undeniable economic advantage, automated tape storage can provide a configuration flexibility unrivaled by other media.

CASE STUDY: Blue Cross and Blue Shield of Western New York

What does Jim Kelly, quarterback of the Buffalo Bills, have to do with Blue Cross and Blue Shield of Western New York, you ask? Easy. A photo of Kelly's face was the first image Blue Cross and Blue Shield scanned in its April 1993 pilot program to test plans for imaging insurance documents. The pilot program grew into a full-fledged commitment to take advantage of imaging technology, and the StorageTek robotic libraries and Iceberg are integral to the success of this mighty adventure.

The Challenges

Blue Cross and Blue Shield recognized the need to improve its paper-based work processes, which impacted its ability to provide timely and accurate service to its customers.

"There are three things that are key to our success for this particular project," said Gary J. Kerl, vice president, corporate information services, Blue Cross and Blue Shield of Western New York. "First is why we installed your silos; second, our imaging system and how it is meeting our goals and objectives; third, why we use the silos for storing our image data."

First, Why Silos?

Simple. The tape library was growing substantially. The number of cartridges was increasing and, from an operational standpoint, it took more people to handle the work while mount errors were, well, mounting. Silos were justified and paid for with offsetting costs in computer operations and better use of technology.

With two PowderHorns installed, Blue Cross and Blue Shield met its automation and quality goals.

Second, Why Imaging?

Here's Kerl again. "We had a corporate need to resolve the paper issues inherent in our work and, through that, to service our customers better. Image processing is key to us moving to a paperless office."

Third, Why Silos?

"Optical was too new and, at the time there were no standards, so we began looking at an alternative for a second level of storage, below our 3390-9 DASD," said Kerl. "We already had the silos installed and wanted to expand the use of that technology, if applicable. It has worked for us."

Blue Cross and Blue Shield has two PowderHorns and uses Silverton 36-track with long tape. They currently have 160 gigabytes of image data online and about 400 gigabytes on tape.

Some Background on the Imaging Decision

Blue Cross and Blue Shield began the imaging pilot program in April 1993 with about 250 claims a day to see how this plan would work. It was a surprisingly smooth transition, although critical to the success of the pilot were two elements: teamwork and significant planning time before implementation.

Now they are scanning about 11,000 medical claim forms and related documents a day; this represents about 5 gigabytes of data images a week.

Blue Cross and Blue Shield's primary goal is to implement EDI for total elimination of paper using electronic data transactions. They currently receive about 75 percent of claims transactions electronically; the remaining 25 percent is paper-based.

Additionally, in order to speedup the keying and processing, Kerl said a substantial amount of paper was being sent out-of-house to an independent contractor for keying. That added time, error, and expense to the process. Another option, microfilm, was messy, slow, and difficult to store and retrieve.

What About Optical for Second-level Storage?

Image data is not normal coded data, that is, it is digitized or object data. "You cannot store that data online forever," said Kerl. "You would use too much storage and it would be too costly. The response time would be good, but response time would in no way pay for the amount of storage you would need."

Implementing imaging meant choosing a second level of storage. According to the folks at Blue Cross and Blue Shield, the normal alternative would be optical disk, but an optical system compatible with the Blue Cross imaging system was too costly. In addition, optical disk storage standards were uncertain; Kerl and the team didn't want to commit to a technology with strong potential for obsolescence.

Why Wasn't Tape a Choice?

Second-level storage was really the second part of a two-part process. The first decision Blue Cross and Blue Shield had to make was whether to use high-performance DASD or some other technology, in particular fat DASD, to hold the initial scanned data. They chose fat DASD.

Mike Morrisoni, manager, systems programming, at Blue Cross and Blue Shield. "Then, after you start to fill up the DASD farm, do you buy more DASD or start to look down the storage hierarchy? When we began our program, optical was our only other choice. There was no application software available to move image data from DASD to tape." That was then.

This is now: IBM ImagePlus software made tape an option. Jeff Kawa, systems programmer at Blue Cross and Blue Shield: "We had two problems—and they got mixed—how do we automate our operations and provide for long-term storage for our data, particularly our image data. With the silos, we killed two birds with one stone."

The Benefits are Legion

In the old system, the claim form journeyed from the mailroom to the computer in 11 or 12 separate steps that took seven to nine days.

In the new system, the journey takes two days. Additionally, once the form is scanned, it is accessible. When the paper is in an electronic queue, the image can be sent around the building within seconds.

In the old system, during these seven to nine days the claim file was "in process"—i.e., not available or cannot be found- to any subscriber who called for information or needed help. Even the Blue Cross and Blue Shield employees didn't always know where in the system a file was.

In the new system, information is available to subscribers, supervisors, and workers very, very quickly.

In the old system, a supervisor who was trying to balance workloads among employees would have to contact each employee in the group, count the number of files each was working on, and redistribute the files manually. In the new system, because the claims are scanned, they are counted automatically and the counts are available on demand to the supervisors. They know their workload and they know it much earlier in the process.

Kerl continued, "There is no doubt that automation represented a decrease in staff costs, handling, and accuracy just within the computer operations department. By implementing the new data center technology, including the silos, my data center hardware budget decreased by 20 percent and our internal competitiveness improved as well.

"And, with the imaging system, we have substantially improved the external competitiveness of the corporation." How so? An industry measurement and regulatory association sets standards for providers and scores member companies on responsiveness, and how quickly claims are settled or customer inquiries handled, for example. Blue Cross and Blue Shield sees improvement here tied directly to the imaging implementation.

Kerl remembered that one of the biggest process improvements Blue Cross and Blue Shield had made in the past was turning microfilm requests around overnight—at best. "Since we have converted to imaging, we can get a retrieval of an imaged document in seconds." The benefit here is far faster responsiveness and better customer service.

In another measure of the benefit of imaging, Kerl said they saved $500,000 annually in data entry expenses by switching to scanning technology for just the medical claim forms. As other paperwork is scanned and put into the electronic workflow process, further savings are expected.

The Next Steps

Right now, according to Morrison and Kawa, they are averaging about 7,000 retrievals per day on the imaged claim forms. Online response time is in tenths of a second; retrieval from tape is, on average, about 40 seconds. The performance is acceptable and the cost is reasonable.

The goal now is to determine in advance which retrievals would be from tape and which from disk. If Blue Cross and Blue Shield can hold likely-to-be-accessed data on disk and move the likely-never-accessed-again data onto tape immediately, they can improve response time to the client and maximize the value of their technology. "This is absolutely a solvable problem," said Kawa. "We tend to think that there is something that statistically marks a claim more likely to be retrieved. We just haven't had the opportunity to study the data enough to understand."

Another goal is to step up the pace for getting many other kinds of documents into the imaging system and/or the electronic workflow. Inquiries and correspondence from subscribers are big targets but, really, scanning any and all paper is targeted, including normal business activities such as Blue Cross and Blue Shield payroll and administrative data.

And What Advice Have You for Others?

"I would tell people that putting imaging on top of current business processes is not going to buy you the gains in productivity and the reductions in cycle time you are looking for." That's Mary Angelo, project manager, systems development. "I would never put imaging in place without reengineering the business process because, typically, you are going to take whatever messy manual process you have and put a piece of technology on top of it—unless you are fortunate enough in your business to have a streamlined and efficient manual process already in place. I don't think there are many of those."

Angelo continued. "I would also look at electronic workflow as part of the solution. In other words, imaging in and of itself isn't enough. Without electronically moving that image around to the right people, you won't get the benefits."

The Bottom Line

Blue Cross and Blue Shield has been able to control costs and still provide a level of service that is acceptable to the end user. They have increased competitiveness through business processing reengineering and the application of imaging technology. They justified the silos to get to automated operations—and got secondary storage for their image data besides.

Glossary

Automated Cartridge System Library Software (ACSLS)—Library control software which runs on a UNIX system to allow client systems to utilize the Automated Cartridge System (ACS). Each client system must have a Client System Component (CSC).

Automated Cartridge System (ACS)—Also known as the Library. Fully automated cartridge storage and retrieval system. Each system consists of: a minimum of one Library Management Unit (LMU), a minimum of one Library Storage Module (LSM), one Library Control Unit (LCU) per LSM and the appropriate Cartridge Tape Subsystem. Multiple LSMs may be attached via Pass-Thru-Ports (PTP) and managed as a single storage entity under the control of one LMU.

Audit—An ACS operation performed to establish a catalog of VOLSER labels, and their physical locations, within the ACS.

Cartridge Access Port (CAP)—Allows cartridges to be entered into, or ejected from, an LSM without opening the main library access door thus maintaining the integrity of the Control Data Set (CDS). All CAP transactions are logged by the ACS.

Computer Assisted Retrieval—An automated document storage and retrieval technology that uses computer hardware and software to index and locate documents or document images recorded on any media.

Cartridge Drive—A device containing a cartridge tape transport and read/write mechanism along with the associated pneumatic systems and power supplies.

Control Data Set—The data set created via the Audit function and used by the host software to control the functions of the ACS.

Client Link—The communications link between the Common Library Server (CLS) and a client. A software component of the storage server which resides in the client processor and interfaces with the CLS.

Client Link Emulator—A PC-based diagnostic tool which monitors a client link and can emulate either the Logical Port or the CSC.

Client System—The system to which the CLS provides an interface to a StorageTek ACS.

Client System Interface—Software which provides a transport and translation mechanism for messages between the Library Control System and the Client System Component (CSC). The Client System Interface provides levels 5 (transport) and 6 (presentation) in the ISO/OSI model.

Computer Output Microform (COM)—A method of mass storage using microfilm or fiche.

Control Unit (CU)—A microprocessor-based unit situated logically between a host channel (or channels) and the tape transport(s). It translates channel commands into transport commands, sends transport status to the channel(s), and passes data between the channel(s) and transport(s).

Control Path Adapter (CPA)—A hardware device which converts from a Client Computing System's control protocol to the control protocol of the StorageTek Library Control System.

Cost per Input/Output (CPIO)—STK proprietary software used to analyze a computer system's I/O traffic in order to determine the appropriate storage configuration.

Client System Component—Software which runs on a client system to interface to the operating system and communicate with a library control server. There are two types of CSC: those that communicate with the

VM Library Control Server (VM protocol) and those that communicate with the UNIX Library Control Server.

Data Path Adapter (DPA)—A hardware device which translates from a Client Computing System's data protocol to the data protocol of the StorageTek Control Unit.

Data Facility Hierarchical Storage Manager (DFHSM)—An integral part of IBM's DFSMS™ (Data Facility Storage Management Subsystem) architecture, it manages low-activity and inactive data in a hierarchy of storage devices having different costs, capacities, and access (performance) attributes.

Enterprise Systems Connection (ESCON)—IBM's fibre optic channel protocol for enhanced connectivity options.

ESCON Converter—ESCON to FIPS 60 protocol converter. Allows attachment of a non-ESCON capable 1/0 device to an ESCON channel.

Expert Library Manager (ExLM)—Optional software which runs under HSC on MVS systems. Facilitates contents management for ACS libraries. Recommends cartridges to be included or excluded from an LSM based on frequency of use, desired number of empty slots, and desired number of scratch tapes.

Host Software Component (HSC)—Resembling an intelligent operator living inside your host system, HSC tracks control data sets (like library catalog cards) and cartridge status (i.e., scratch) inside the ACS. HSC intercepts regular operator mount/dismount messages and communicates with the LMU for processing by the LSM (robotics).

Improved Cartridge Recording Capability (ICRC)—StorageTek's term for data compression and data compaction on 18-and 36-track cartridges. ICRC is 100% compatible with IBM's Improved Data Recording Capability (IDRC).

Installation Verification Programs (IVP)—A package of programs that is run by a user after the library is installed in order to verify that the library is functioning properly.

Library Communication Facility (LCF)—Software which provides the required TCP/IP support.

Library Control Unit (LCU)—The portion of an LSM that controls the robot's movements as directed by the LMU. The LCU is integrated into the smaller LSMs but is an additional unit attached to the side of the large LSMs.

LIBGEN—The process of defining the configuration of a library to the host software component.

Library—A library is composed of one or more ACSs, host software that controls and manages the ACSs, and the library control data set that describes the state of the ACSs.

Library Control Component (LCC)—Software which controls the mounting and dismounting of cartridges in an ACS.

Library Control Software (LCS)—Library Control Component + Client System Interface + Library Utilities.

Library Station—Software which allows an MVS/XA or MVS/ESA HSC to be utilized from other (client) systems as a server. This vehicle also allows other non-MVS based subsystems to share as library with existing MVS/HSC based systems as well a providing the basis for a better solution for remote shared libraries.

Library Management Unit (LMU)—The portion of an ACS that controls the LSM and communicates with the host software component.

Library Storage Module (LSM)—A housing containing cartridges and a robot that moves the cartridges between their storage cells and the attached cartridge tape transports.

LSM Number—A hexadecimal value from O to F.

Maintenance Facility—Hardware in the CU and in the LMU that allows a field engineer and the Remote Diagnostic Center to run diagnostics, retrieve status, and communicate with these units through their respective control panels.

NearArchive—A tape database software product which stores, manages, and indexes data to tape thus enabling applications to read and write data directly to and from tape.

Nearline™—StorageTek's architecture for a level of storage positioned between online and offline, that improves the performance of tape access in a cost effective manner.

NearOAM—A software module used in conjunction with NearArchive which enables any OAM (IBM's Object Access Method) application to sort, manage, and access data from tape.

Optical Character Recognition Label (OCR Label)—An external label attached to a tape cartridge that can be read by both humans and machines.

Physical Port—The communications hardware required to support a CLS/Client Link.

PM2—Mainframe software, (PMPMS) Performance Measurement and Predictive Maintenance System.

Program for Online System Testing (POST)—A program in the host that allows it to test an attached subsystem while the subsystem is online.

Pass-thru-Port (PTP)—A mechanism that allows a cartridge to be passed from one LSM to another in an ACS with more than one LSM.

Physical Volume Manager (PVM)—Pass thru facility for Virtual Machine.

Remote Electronic Storage Solutions (RESS)—AKA electronic vaulting. RESS allows a data center to write its backup tapes in a location geographically separate from the production data center thus avoiding the need to ship tapes offsite.

Shadow Recording—A technique for recovery that maintains two identical copies of the Control Data Set.

Transport—An electro-mechanical device that threads tape from a cartridge past a read/write head to facilitate data transfer. Located in the CD portion of a tape subsystem.

ViewDirect—A software product that offers direct, online access to report archives across a hierarchy of storage media. ViewDirect matches a client's report access frequency, the speed of available storage devices, and the cost of storage to optimize long term retention of report archives. Co-developed by Mobius Management Systems and StorageTek, ViewDirect presents a practical, reliable Computer Output Microform (COM) alternative.

VIP—StorageTek customer.

Volume/Serial Number Label (VOLSER Label)—A six or seven character alphanumeric label, often with accompanying bar-code that identifies a tape volume.

Volume—A data carrier that is mounted or dismounted as a unit.

Write Tape Mark (WTM)—An operation that records a special mark on a tape thus identifying a specific location on the tape.

NOTES

NOTES

Visit ATG's Web Site
to read, download, and print
all the Technology Guides
in this series.

www.techguide.com

"The significant problems we face cannot be solved
by the same level of thinking that created them."

Albert Einstein

RHEINNER GUIDE

This Technology Guide is one of a series of guides, written by The Rheinner Group and published by ATG, designed to put complex document management, imaging, and workflow concepts into practical and understandable terms. Each guide provides objective, non-biased information to assist in the internal education, evaluation and decision making process. This Technology Guide, as well as the other Document Management, Imaging, and Workflow Technology Guides in the series, are available on ATG's Web Site.

www.techguide.com

A joint venture

ATG
The Applied Technologies Group

Tel: (508) 651-1155
Fax: (508) 651-1171
E-mail: info@techguide.com

The Rheinner Group

Tel: (617) 741-8100
Fax: (617) 741-5885
E-mail: info@rheinner.com