

DIGITAL ARCHIVING SOLUTIONS FOR THE MEDIA AND ENTERTAINMENT INDUSTRY

White Paper
February, 2008

Table of Contents

Challenges of long-term, media content archiving	1
Benefits of digital archiving	2
Overview of Sun’s digital archive solutions	2
Digital tape — the safe choice	4
Key considerations on implementing a digital archive	5
Capacity	6
Virtually unlimited near-line capacity	7
Online capacity for production	8
Managing massive digital media archives	8
Performance	10
Retrieve content faster	10
Software increases archive performance	11
Increase operational performance with integrated array	12
Reliability and availability.	13
Engineered for data intensive environments	13
Designed for data integrity	13
Tape drive reliability	14
Tape system availability	14
Cost of ownership	16
Minimizing capital expenses	17
Controlling media costs	17
Reducing facility and operating costs	18
Conclusion	19
References	20
Web Sites	20
Papers	20

Chapter 1

Challenges of long-term, media content archiving

Since the early years of the entertainment industry, film makers have struggled with the problem of storing motion pictures for long periods. Thousands of silent films were made in the years leading to the introduction of sound, and between 80 and 90 percent have been lost forever — many because of the deterioration of the actual celluloid film. Movies from the first half of the 20th century were captured on an unstable, highly flammable, cellulose nitrate film base, which required careful storage to slow its inevitable process of decomposition. Most films made on nitrate stock were not properly preserved and, over the years, their negatives and prints simply crumbled into dust. Many early celluloid films are damaged by dirt, scratches, tears, color fade, color change, film grain noise, not to mention missing scenes and sound that might have been censored or edited out.

Similar difficulties occur during the long-term storage of video programming. Video tapes consist of an acetate, polyester, or mylar base coated with an oxide. The oxide is covered with a lubricant to help minimize friction as the tape passes over the read/write heads of video recording and playback machines. Over time, the lubricant evaporates or undergoes a chemical change rendering it useless. Once the lubricant is gone, flakes of oxide containing recorded information wear off each time the tape is read. Valuable information is lost and the tape heads become clogged. Depending on the composition of the substrate and the binder (glue that holds the oxide to the base), the tape might develop clear spots where no oxide remains resulting in the loss of recorded material. Also, the substrate can become brittle making it too delicate to pass through a video playback machine without breaking apart.

The extensive — and expensive — process of accessing and reusing media is an additional major problem with traditional, analog archives. When a celluloid film is retrieved from an archive, it must be located, pulled from the vault where it is stored, brought up to ambient temperature and humidity, packaged, insured, and shipped to the requesting party. Once the archived film is delivered to the requestor, it must be inspected to determine if it needs to be restored before it can be reused. This process can take up to two months and cost as much as \$10,000.

In addition, much of today's programming is created, produced, and distributed digitally. Consequently, to reuse an archived celluloid film or analog video in a new production, it must first be scanned and digitized — another time consuming and expensive process — after which it can be digitally restored and integrated into the new content.

Finally, the mundane task of manually handling video tapes or films can be both time consuming and space consuming as well as fraught with the risk of losing or damaging the content. Media can be mislabeled and misfiled making it virtually impossible to locate in a large vault. A video tape can be dropped or carried near a strong magnetic field rendering it worthless. Celluloid film can be put through damaging heat shock when it is moved from a hot delivery truck into an air conditioned studio, irreversibly corrupting the images.

Benefits of digital archiving

Virtually all of these problems are eliminated by implementing a digital content archive. And, since digital media is replacing analog video tapes and celluloid film in all phases of the imaging production, many producers believe the digital process should be embraced fully by the industry now that the technology exists to do so cost-effectively. Digital content can be stored electronically near or on the production system, making the content readily available for editing or reuse without the time, cost, and risk inherent in locating it in a vault and physically transporting it to a production facility. Nor is it necessary to undergo the time-consuming and expensive film-out process to convert celluloid film to a digital intermediate (DI). Digital films and video programs are readily and cost-effectively accessible for reediting and reuse. As a result, many media houses are considering upgrading their production facilities to include digital archives.

Digital archives help ensure that content placed there today is available for viewing and reuse well into the future. The cinema industry requires that films are stored and available for about 100 years and the video industry is storing programs for 50 years and more. Through the prudent use of best archival practices, including periodically upgrading hardware and software technology and making copies of digital content at appropriate intervals, digital media archives can be expected to securely store content for 100 years and more. (See *Archiving Movies in a Digital World*, a white paper by Dave Cavena et al, January 10, 2007, for a description of the major aspects to implementing and operating a digital archive for film).

Overview of Sun's digital archive solutions

Sun Microsystems offers the servers, automated tape libraries, disk storage, management software, and infrastructure for creating and maintaining a digital media archive. As shown in Figure 1, the heart of the digital media archive consists of a high-capacity, automated tape library providing the necessary long-term storage, integrated with a high-performance disk system delivering digitized content for online editing and post-production. The backbone of the system is a Fibre Channel (FC) interconnect, providing the Storage Area Network (SAN) over which content files are stored and shared, and a high-speed Gigabit Ethernet for system communications. Sun™ servers running shared file system software and hierarchical storage management software control the environment.

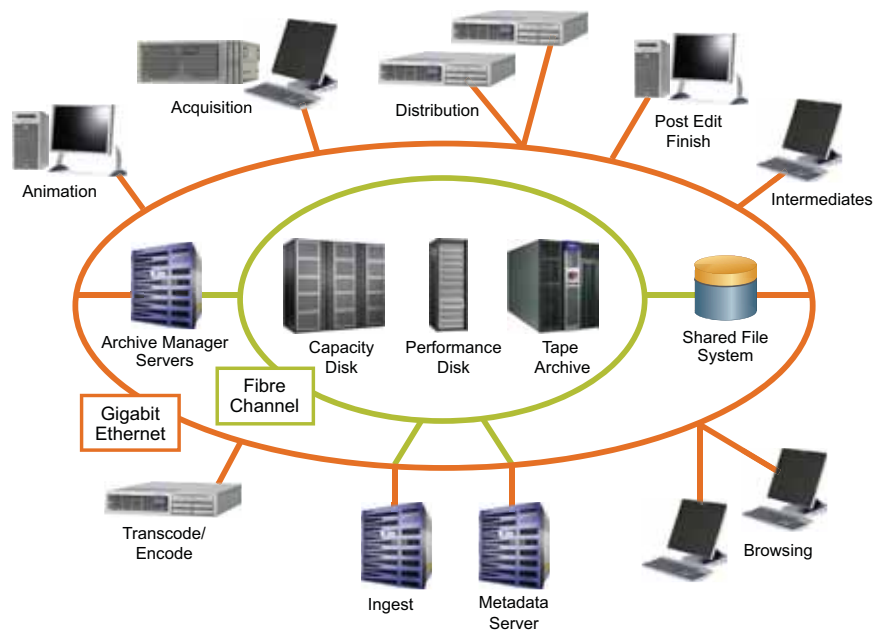


Figure 1. A typical configuration for a midrange digital media content archive.

A typical archive system configuration includes:

- Long-term tape storage — the Sun StorageTek™ SL8500 modular library system, an extremely scalable and flexible tape library, is ideal for media archives. Using Sun StorageTek T10000 tape drives, the capacity of the library scales from 724 terabytes (TB) up to 35 petabytes (PB) and supports from 64 to 448 drives. For environments that do not need these high levels of capacity and throughput, the Sun StorageTek SL3000 modular library system scales from 100 TB to 1.5 PB of tape storage, supports from 1 to 56 tape drives, and offers the same true mixed media, partitioning, advanced management, and redundancy features of the Sun StorageTek SL8500 tape library.

- High-capacity tape storage — the Sun StorageTek T10000 drive is a large-capacity (500 GB per cartridge) and high-throughput (120 MB per second) tape drive that is designed to be extremely reliable in the high duty-cycle, media production environment.
- High-performance 1st-tier storage — the Sun StorageTek 6540 array is a scalable, highly reliable disk array designed to support up to 224 disk drives for a total of 168 TB of storage, providing sufficient online capacity and performance for most production operations.
- Cost-effective production processing — a wide range of highly-scalable SPARC®, Intel® Xeon®, and AMD Opteron™ processor-based servers that cost-effectively and reliably meet the performance demands of virtually any media production house.
- Powerful storage management — The combination of Sun StorageTek QFS and SAM-FS software is a tightly integrated data management solution designed to optimize storage hardware, providing innovative file system, file sharing, and archiving services.

Even though other hardware and software combinations might support faster data transfer rates or higher storage capacity, this paper focuses on this suite of components since it offers the best overall capabilities for cost-effectively archiving and managing digital film and video content.

Digital tape — the safe choice

Media production companies have a number of choices for storing data, including optical disks, magnetic disks, CD-ROMs, and DVDs. Yet digital tape, the computing industry's first reliable format for long-term storage, continues to play a major role in the storage strategy for digital media firms. In fact, many companies still view tape storage as their *last line of defense* — a fail-safe location where information is preserved when it exists nowhere else.

Tape is the most affordable storage solution available. Tape is portable so it can be easily stored in a safe and secure facility. Tape is dense — a single Sun StorageTek T10000 tape cartridge can hold 500 GB of data, which is ten times the capacity of a standard Blu-ray optical disk. Tape is high performance — the Sun StorageTek T10000 tape drive is capable of transferring data at 120 MB per second, more than 15 times Blu-ray's data rate. But, most importantly, tape is safe. It uses smart automation software and backup and recovery processes that are validated worldwide. Tape's automation and virtualization capabilities are legendary, trusted by some of the world's most successful companies. And now with device-level encryption, a lost or stolen tape that has been recorded on the Sun StorageTek T10000 (Figure 2) can be rendered useless to anyone who tries to read it without authorization.



Figure 2. The Sun StorageTek T10000 Tape Drive is a high-performance, cost-effective tape storage system

Sun continues to drive tape technology forward, developing libraries, drives, and media that safeguard data, accelerate access, and reduce the cost of storing and moving data. Sun tape libraries and drives are easy to install, scale, manage, and maintain. The Sun StorageTek tape portfolio offers a broad range of capacities and performance levels to meet the operational and budgetary requirements of virtually any production house.

Tape drives and tape media available from Sun offer unprecedented durability and dependability. Highly accurate error correcting code (ECC) helps ensure data integrity while improved tape guiding systems designed to reduce tape speed with fewer tape passes help to reduce drive and media wear. Sun also offers industry leading tape management software to help increase the efficiency and productivity of the automated tape libraries with streamlined management, better monitoring of backup applications, and user-friendly interfaces.

Key considerations on implementing a digital archive

A digital archive is generally considered to be much more than simply a passive repository for data. It also embodies the technology and processes that access, manage, and distribute content to editors and production professionals for additional repurposing opportunities. This paper addresses each of the four key functional areas that must be considered when contemplating the development of a digital archive system:

- Capacity
- Performance
- Reliability and availability
- Cost of ownership

Chapter 2

Capacity

Thought Equity Management, Inc.

Thought Equity is a leader in motion imagery licensing and management for the entertainment, creative, and corporate production industries. Creative professionals get instant access to the world's finest film and high definition content on Thought Equity's Web site.

The company's business was growing dramatically and they needed a content management system that could cost-effectively scale to petabytes of capacity while providing customers with easy and rapid access to their film and video files.

After adding a Sun StorageTek T10000 tape drive to its existing tape library, the storage density, speed, and scalability of Thought Equity's digital archiving and on-demand content delivery system was dramatically improved.

The company can now deliver content four times faster while storing two and a half times more data in the same footprint. Furthermore, their petabyte-level storage environment scales in cost-effective increments to keep pace with their expanding content archive and to meet the growing demand for their services.

The amount of film and video content recorded in the media and entertainment industry is growing exponentially, and today, much of this content is recorded digitally. All-digital motion picture and video production is becoming increasingly common due to the advent of new digital camera technology that is suitable for shooting feature films. Capturing images directly in digital simplifies post-production work such as editing, sound, color correction, and special effects. And because digital content does not consume film nor does it require pre-processing to create digitally editable content, directors are allowing cameras to run much longer while shooting scenes, increasing the opportunity to capture unique moments that might otherwise be lost.

At its most basic level, the issue of capacity is quite simple — acquire enough storage for all of the digital content. Typically, an hour of high definition (HD) video requires about 23 GB of storage and a digital intermediate (DI) version of a full-length feature film consumes about 10 TB.¹ By determining the quantity of movies, videos, and other programming to be stored, an estimate of the size of the archive can easily be made. But, there are additional aspects of capacity that need to be considered including growth of the archive and accessibility of the content.

Storage space sufficient to accommodate productions that are expected to be worked on in the foreseeable future must be made available. It is generally unwise, however, to project the size requirements of an archive beyond two or three years since the assumptions on which the sizing is based are unlikely to be valid beyond that time. Additionally, due to changes in storage technology, what appears to be more than sufficient storage now might be inadequate in five years. Regardless, since unexpected and sometimes rapid content growth is almost inevitable, the design of the archive should be scalable in easy and cost-effective increments.

Once the total amount of storage is determined, the type of storage — online, near-line, and offline — needs to be considered. Online storage is for content that must be available to the production staff immediately and is stored in high-performance, online disk arrays. Near-line storage is for content needed within just a few minutes and is stored in a more cost-effective tape library. This library must be highly scalable since it is the repository that ultimately must accommodate the entire archive and must be able to easily grow to accept unexpected additions. Finally, off-line storage is the physical space needed for the archive vault. It is for content that is not expected to be needed in the near-term and can be loaded onto tapes and stored securely off-site.

1. The digital intermediate (DI) process essentially converts a movie from celluloid film to digital bits for post-production work and then back to film for release to theatres. While the size of a DI depends on the length of the film, for the purposes of this paper a DI is assumed to be 10 terabytes (TB).

Virtually unlimited near-line capacity

The Sun StorageTek SL8500 modular library system (Figure 3) is uniquely designed to provide virtually limitless storage capacity while allowing growth in easy-to-manage, cost-effective increments. A basic Sun StorageTek SL8500 library can accommodate up to 1448 cartridges providing 724 TB (uncompressed) of near-line storage capacity offering users access to as many as 72 feature film DIs. (See footnote on page 7.)



Figure 3. The Sun StorageTek SL8500 library system enables extremely scalable and space-efficient tape archives

Each Sun StorageTek SL8500 library is equipped with four (optionally eight) HandBot robotics working in parallel to handle multiple, asynchronous mount requests. Up to seven library modules can be interconnected and operated as one system. Since each additional library module contains its own four (or eight) Handbot robotics, the performance does not degrade as capacity is added — as it does with many traditional tape libraries.

The Sun StorageTek SL8500 library includes a minimum of 1448 cartridge slots and can grow to 10,000 cartridges — for 5 PB or 500 feature films. (See footnote on page 7). Multiple Sun StorageTek SL8500 modular tape systems can be attached via pass-through ports creating a single logical library that can accommodate 448 tape drives accessing 70,000 cartridges providing up to 35 PB of storage — enough to provide rapid access to as many as 3500 feature films.

Based on the innovative design of the Sun StorageTek SL8500, the mid-range Sun StorageTek SL3000 modular library system supports up to 56 tape drives and offers non-disruptive scalability up to 3000 slots — providing as much as 1.5 PB of storage capacity when used with Sun StorageTek T10000 tape drives and cartridges. Of course, with either library system, additional cartridges containing even more films can be vaulted off-site for later use.

Online capacity for production

The Sun StorageTek 6540 disk array stores up to 168 TB of digitized content (enough for 16 movie DIs) online and available for production work. The array can be installed with just 5 disks and then expanded as needed up to 224 disks with a mix of high-performance FC drives and high-capacity Serial Advanced Technology Attachment (SATA) drives in easy to manage, cost-effective increments. And, as noted in Chapter 4 *Reliability and Availability* in this paper, all of these tape and disk capacity upgrades can be performed nondisruptively — while the system continues to operate.

Managing massive digital media archives

Once the movies and videos are digitized and written into disk arrays and tape libraries, the task of managing and protecting these assets becomes paramount. Sun partners with leading storage management software providers such as EMC, Front Porch Digital, Masstech, and SGL to offer secure, scalable archive systems designed to handle a wide range of media content storage environments. Regardless of the medium or the ultimate production piece, Sun and Sun storage management partners help media and entertainment companies easily and cost-effectively store, edit, reuse, distribute, and archive their digital media content.

For many media production houses, demand is growing for digital archive systems that can support multiterabyte files contained in file systems of tens or even hundreds of petabytes. These capacity requirements strain the limits of traditional file systems. The advanced file system that forms the foundation of Sun StorageTek QFS and SAM-FS storage management software is designed to scale as data storage needs grow, supporting virtually unlimited capacity. Furthermore, the file system used by Sun StorageTek QFS and SAM-FS software manages the growth of storage capacity automatically, reducing the need for risky and labor-intensive reallocation of storage into and out of existing file systems as their sizes change. The storage system's cache is

sized for the expected working set of files, rather than the required total capacity. User-defined production rules describe the relative importance of content and the system manages the data appropriately, moving it to near-line and then to offline storage when its importance diminishes. And all of this happens on a high-performance file system that can operate and be shared at device speeds.

Consolidate, share, and protect media content

Sun StorageTek QFS software is a highly scalable, shared file system, designed to provide file sharing over SAN, where all data is consolidated into a pool and is accessible to multiple users and applications. It is designed to provide maximum scalability, performance, and throughput for data-intensive applications such as digital media creation without the need to maintain and synchronize different file systems for different production purposes. With Sun StorageTek QFS software, massive data sets do not need to be spread across multiple file systems — they can be stored in a single, scalable file system, significantly reducing administrative overhead without degrading performance. Multiple applications and users can share the same files and volumes on the network, dramatically improving productivity of the content production process.

Archiving content cost-effectively

Sun StorageTek SAM-FS software is a Hierarchical Storage Management (HSM) solution that helps improve the storage capacity utilization of the media production system as well as the productivity of the production process by enabling dynamic archiving, reduced backup windows, and fast restore. Administrators can set automatic archiving policies that determine when, where, and how information is stored, cost-effectively managing content based on its production value. For security and performance, up to four copies of data can be archived simultaneously to local or remote disk, tape, or optical storage.

Chapter 3

Performance

Traditionally, obtaining video or film content requires locating the item in a library vault, shipping it to the user, and manually inserting it into an available playback or scanning device — a process that can sometimes take months. Today, a digital archive user expects the system to deliver a stored program or movie to their desktop in minutes. The primary features that help enable the archive to meet the performance these users demand are:

- High data rate at which content is downloaded from the near-line tape library into the production system
- Simultaneous data transfers from multiple tape drives increasing the effective throughput of the tape library
- Sophisticated file system software helping to optimize the handling of large data sets
- Use of online, high-performance disk storage to cache current production files for immediate access

Retrieve content faster

Although much quicker than obtaining a celluloid film or video tape, retrieving a single stream of digitized content can also take a long time. Reading a 10 TB DI file on a Sun StorageTek T10000 tape drive at 120 MB per second takes approximately 23 hours not including the time to load and unload the 20 or so tape cartridges. Since, however, the basic Sun StorageTek SL8500 library supports as many as 64 Sun StorageTek T10000 tape drives, the DI retrieval process can be multithreaded, enabling all 20 cartridges to be read simultaneously. Concurrently reading 20 tapes effectively provides a data transfer rate of 27.6 TB per hour reducing the time to obtain the movie DI to approximately one hour.

Simultaneous reading of multiple data streams from multiple storage devices is supported by the advanced file system used in Sun StorageTek QFS and SAM-FS software helping production facilities access their content faster and complete their projects on schedule. Since both the Solaris™ Operating System (OS) and the advanced file system support multithreading, multiple threads can access multiple tape drives at the same time, a capability that is extremely beneficial for applications such as video streaming. Sun StorageTek SAM-FS software also allows each copy or restore operation to function as a fully independent thread, enabling the software to stream data to or from each tape drive at its rated device speed.

Software increases archive performance

With unique features for handling extremely large files such as video clips and digitized motion pictures, Sun StorageTek SAM-FS and QFS software help enhance the performance of a digital archive. The underlying file system used by the software is designed to minimize the number of instructions required for each read/write transaction and to intelligently optimize data layout and input/output (I/O) patterns on the storage devices. These features help to maximize the speed at which media content is transferred between the tape library and the production system helping the production staff to complete their projects sooner.

Sun StorageTek QFS software helps reduce performance bottlenecks by optimizing the operation of the file system in conjunction with the underlying disk technology and hardware. It helps overcome other UNIX® file system shortcomings such as the time required to create new file systems, file system growth limitations, lengthy file system checks after an unintended interruption, and limitations on the number of files in the file system. To increase throughput, Sun StorageTek QFS software minimizes the movement of the physical head on the disk with new technologies such as metadata¹ separation, variable block size, pre-allocation, and automatic direct I/O. Also, integrated volume management with performance options such as disk striping and disk allocation eliminates additional software layers and uses less system overhead than other UNIX file systems, improving the overall productivity of the archive.

Sun StorageTek SAM-FS software consolidates innovative archiving and backup methodologies into its high-performance file system. The software replaces traditional backups to improve storage resource utilization for applications where data needs to be available continuously and quickly restored in the event of a business disruption. Metadata archiving and read-behind features help production facilities recover from disruptions in minutes or hours, as opposed to days, and let users begin reading files even before they are fully restored.

The Solaris OS helps Sun servers achieve superior performance and scalability by using threads² as the fundamental technology employed to allocate processors. Many threads can be executed in parallel by separate processors in multiprocessor systems. The Solaris OS responds to interrupts, performs driver and background activities, and handles application requests using threads, and the number of possible threads is limited only by the amount of available memory.

1. Metadata provides information to the file system about the characteristics of the data stored.

2. A thread is an independent sequence of program instructions that can be executed by a processor.

By using the same mechanism for scheduling both system activities and user processes, improvements in threads improve all aspects of performance for applications built to utilize multiple processors, including the advanced file system in Sun StorageTek QFS and SAM-FS software. In an environment where both the operating system and file system support multithreading, multiple threads can access a file at the same time, which is beneficial for applications such as media streaming and large media file production.

Increase operational performance with integrated array

Due to the relatively lengthy operation necessary to obtain a digital file directly from a tape library¹, a tape-only archive system does not have sufficient performance to operate in a production environment. To resolve this, a high-performance disk storage system, such as the Sun StorageTek 6540 array (Figure 4), is coupled with the tape library so that digital content is read from the tape library and cached in the disk array enabling post-production functions to be performed in real time.



Figure 4. Sun StorageTek 6540 Array is designed to handle large data sets and delivers industry-leading performance.

1. For example, an entire 10 TB film DI takes approximately one hour to be downloaded from the tape library into the production system using 20 Sun StorageTek T10000 tape drives operating in parallel.

Chapter 4

Reliability and availability

Electrofilm USA GMBH

Electrofilm provides digital media services — DVD production, audio and video restoration, content delivery, and project management — to the media and entertainment industry. The company needed a digital asset management system that included primary and secondary disk storage, tape storage, and management software enabling them to scale easily and to control the entire data lifecycle with a minimum amount of human intervention.

Electrofilm chose a system from Sun that includes Sun StorageTek 6140 and Sun StorageTek 6540 disk arrays for primary and secondary disk storage, a Sun StorageTek SL8500 tape library with Sun StorageTek T10000 tape drives for archival storage, and Sun StorageTek SAM-FS and QFS storage management software running on a Sun Fire™ X4600 server for content lifecycle management.

Using Sun's server and storage products, Electrofilm has created a post-production system — based on a tapeless workflow — in which there is no analog video tape anywhere in the process. This state-of-the-art system enables the company to lower the total cost of the archive by using tiered storage and content lifecycle management. Electrofilm has also streamlined its workflow operations by eliminating video tape and can more effectively compete with other companies that use a mix of analog and digital.

One of the major concerns of the industry with digital archiving is ensuring the production process is able to continue to operate in the event of a hardware or software disruption and the data captured in the production process remains unchanged once it is stored in the archive. It is essential that equipment used in a digital media archive utilizes state-of-the-art error correcting code (ECC) technology to help ensure it precisely reproduces the data. In addition, storage drives and arrays must be designed with the necessary high-availability technology to help ensure content production continues even when a component fails.

Engineered for data intensive environments

The Sun StorageTek T10000 tape drive and tape cartridge are designed to work in high duty-cycle¹ environments — operating 70 to 100 percent of the time. Both the drive and the cartridge include advanced technologies that minimize wear resulting in higher reliability. The Sun StorageTek T10000 drive uses a unique, 32-channel, dual-head technology — enabling slower tape speeds, fewer passes to fill a cartridge, and better utilization of the tape media — which helps ensure reliable reading and writing. Slower tape speed — 4.95 meters per second versus 6.2 meters per second for a typical datacenter-class tape drive — helps minimize the wear and tear on the media and the drive while maintaining high capacity and high performance rates.

Designed for data integrity

When writing bits to a computer tape, the Sun StorageTek T10000 tape drive uses a sophisticated dual ECC scheme and a cyclic redundancy check (CRC) to provide one of the highest levels of error detection and correction available. This helps ensure that if the media is damaged or information on the media is compromised, the drive can detect it and can reconstruct the data from the ECC bits recorded on the tape. The ECC scheme for the Sun StorageTek T10000 tape drive has an uncorrected bit error rate of 10^{-19} resulting in a data accuracy rate two orders of magnitude better than that of competitive tape drives. In addition, the dual head design enables the data to be spread over a larger area on the tape making the drive less susceptible to data drop outs and tape defects.

Furthermore, the Sun StorageTek QFS software maintains the integrity of the file system with features such as integrated error checking on all critical I/O, serializing critical metadata writes, and recording identification records on metadata, which can

1. Duty-cycle is the proportion of time that the device is in operation.

be dynamically detected and recovered. These features help ensure that all completed transactions are recoverable.

Sun StorageTek SAM-FS software includes a system error facility (SEF) reporting mechanism, which helps notify system administrators if tapes have aged and should be replaced or rotated. In addition, Sun StorageTek SAM-FS software supports TapeAlert, which provides detailed diagnostic information using an industry standard interface. A tape drive that supports the TapeAlert standard constantly checks the drive for potential failures and reports any problems before they occur.

Even the Sun StorageTek T10000 tape cartridge is designed to protect the integrity of the content. If a typical cartridge is dropped, the reel flange can damage the edge of the tape making the tape unreadable. To protect the media and the stored data, the Sun StorageTek T10000 cartridge has a unique, patent pending locking mechanism that prevents the hub from hitting the inside of the cartridge and damaging the tape during rough handling.

Tape drive reliability

The Sun StorageTek SafeGuide system, the tape guiding system of the Sun StorageTek T10000 drive, further protects the media and the data. The unique buckler mechanism securely attaches the cartridge leader to the drive leader protecting the media from damage due to incorrect loading into the drive. The long tape path guides the media more accurately and reduces tape tension reducing wear on the heads and media. The tape guides contact only the back side of the media not the recording surface, minimizing lateral tape motion and protecting data integrity.

The Sun StorageTek T10000 tape drive utilizes a 32-channel, dual-head design, which reduces the number of tape passes necessary to read and write, thus extending head and media life. In addition, the dual head architecture enables the data to be spread out over a larger area on the tape providing increased immunity from data loss. The Sun StorageTek T10000 tape drive has a 256 MB buffer, which, along with its sophisticated buffer management and speed-matching design, enables the drive to achieve optimum performance regardless of the speed of the data feed from the server.

Tape system availability

The Sun StorageTek SL8500 tape library is designed to minimize both scheduled and unscheduled maintenance. All key components in the library, including the tape drives and robotics, are hot swappable enabling them to be removed and replaced without stopping the library's operation. In addition, the power supplies are redundant enabling the library to continue to operate while a failed power supply is replaced.

And with the optional redundant robotics and integrated service safety door, the Sun StorageTek SL8500 tape library can continue to operate even while a failed robot is replaced. In fact, with the optional 2N power configuration, the Sun StorageTek SL8500 tape system continues to operate even if one of the two main electrical services to the library is lost.

The *RealTime Growth* capability designed into both the Sun StorageTek SL8500 and the Sun StorageTek SL3000 enables tape drives and cartridge slots to be activated while the library system continues to operate, minimizing scheduled maintenance downtime. As a result, the library's capacity and throughput can be grown while the system continues to operate. The *Any Cartridge Any Slot* technology provides seamless support for mixed media, helping ensure that any combination of approved drives and media can be installed in the library. Furthermore, there is never a need to take an outage to replace slots to accommodate a different mix of media. Both the Sun StorageTek SL8500 and SL3000 library systems allow new drive and media units to be added without having to install unique slots to accommodate the different technologies, and these performance and capacity upgrades can be performed while the system is operating.

Chapter 5

Cost of ownership

The United States Library of Congress

The United States Library of Congress is the nation's oldest federal cultural institution and serves as the research arm of the U.S. Congress. It is also the largest library in the world, with more than 130 million items on approximately 530 miles of shelves. The collections include more than 29 million books and other printed materials, 2.7 million recordings, 12 million photographs, 4.8 million maps and 58 million manuscripts.

The newly constructed National Audio Visual Conservation Center houses the library's collection of audio and video reproductions. The center is converting millions of sound recordings, videos, film clips, and photographs into digital form in order to protect them from natural deterioration and to improve the public's access to these artifacts.

The center designed an archival infrastructure based on Sun servers and storage. The system includes high-performance Sun Fire X4600 servers running the Solaris 10 OS, a robust storage area network based on Sun StorageTek 6540 arrays, a Sun StorageTek SL8500 tape library with Sun StorageTek T10000 tape drives, managed with Sun StorageTek QFS and SAM-FS software.

With this new media archive, the library has significantly expanded the amount of content it can store and increased the rate at which new content can be acquired. It has also extended the length of time the content can be preserved. As a result, the public has easier and faster access to many, long forgotten audiovisual treasures.

When evaluating whether to deploy a digital media archive or maintain a film or video tape archive, media production houses must consider not only costs but also access to and viability of the associated technologies and services. Many professionals in the media and entertainment industry believe that as shooting and projection join post-production as predominantly digital-based activities, the cost of film and film processing will dramatically increase, if they remain available at all. Consequently, film companies must consider the long-term availability of celluloid film — particularly at a cost that is competitive with digital media. Also the cost and availability of laboratories that can provide the necessary film processing services must be evaluated.

For video production companies, the obsolescence of video playback devices that are compatible with the companies' archived tapes is a major issue. Newer video tape technology usually requires re-recording the content onto new, compatible tapes with the resulting deterioration of video quality, since the copy has all of the noise from the original as well as additional noise in the copy. And, as with film production, all-digital video production is accelerating the end of analog video recording and playback technology.

Digital archives are expected to be operational and provide continuous access to media assets for 50 to 100 years. To meet this requirement, an archive's infrastructure — servers, storage devices, and networking — and media need to be upgraded periodically to make sure the stored content continues to be accessible and readable by the system's devices. Upgraded storage devices typically offer higher performance and capacity in smaller, more cost-effective packages that help to increase the capabilities of the archive while reducing the cost of floorspace, power, and cooling. Archived content should be re-recorded to new tape media periodically to help confirm the tape cartridges and drives are compatible so there is never a 100 year old cartridge with no tape drive that can read it or a 100 year old tape drive with no cartridges it can read. This not only helps prevent problems caused by media deterioration but provides the opportunity to transcode assets to current, state-of-the-industry media and formats, which helps ensure that the assets can be retrieved with the currently available systems. Also, periodically rewriting the content helps keep the tape software updated over time helping to ensure there is never an old tape file without the tape software to access it.

Disk arrays and compute servers should be replaced every five to ten years. Tape drives should be replaced in three- to five-year intervals and tape libraries in ten- to twenty-year intervals. And the published media life for the Sun StorageTek T10000 tape

cartridge is 30 years but should be replaced every ten to fifteen years to help ensure its integrity. Consequently, the total cost of implementing an archive system includes:

- Purchase of the initial system
- Incremental cost to expand the system with additional cartridge capacity, tape drives, and disk arrays as the archive grows
- Operating and maintenance costs of both hardware and software including licenses
- Upgrading and replacing aging equipment on a periodic basis

Minimizing capital expenses

The Sun StorageTek SAM-FS software is a hierarchical management system that determines where content is stored based on user-defined policies. Media content is distributed across multiple levels of storage such as online disk arrays and near-line tape libraries as users define its value to the production process. By storing only active production content on the more expensive, high-performance, online disks, and the less urgently needed content on the more economical, near-line, tape library, the Sun StorageTek digital media archive can provide the accessibility and performance to meet users' production requirements while helping to keep costs to a minimum by reducing the need for additional, more-costly, online storage devices.

The Sun StorageTek T10000 tape drive and Sun StorageTek SL8500 tape library are designed for high duty-cycle environments — in many cases reading and writing in excess of 20 hours per day, 365 days a year. Not only does their high duty-cycle design help ensure the reliability of the system, it also helps to lower equipment costs. For example, a tape system that operates for 20 hours every day must have a duty-cycle rating of 83 percent ($20/24 \times 100$ percent) or more. If the device has a rated duty-cycle of just 50 percent, it can safely operate for only 12 hours each day. The production house would need to purchase an additional tape drive to continue reading and writing tape files for the other eight hours. Thus, high duty-cycle systems from Sun help reduce costs by eliminating the necessity to purchase additional lower duty-cycle tape drives needed to support the throughput requirements of large media files.


Controlling media costs

The Sun StorageTek T10000 tape cartridges can be used for at least two generations of tape drives. Unlike typical competing tape drives that are read-only backwards compatible to the previous generation, the Sun StorageTek T10000 tape drives are designed to be fully compatible, saving operating expenses by enabling the cartridges to be reused at higher-capacity with the next-generation drives. Since the media is often the largest expense in an archive, reusing media for a second generation at a higher capacity can be a huge cost savings.


To reduce the amount of media required, Sun StorageTek SAM-FS software virtually never needs to perform a full backup of all data. Instead, it copies only new and changed files. This saves costs by eliminating the need to add media for multiple copies of old data and saves time and money by reducing administrative overhead to manage all the backups.

Reducing facility and operating costs


The raised floor space needed for a digital archive is expensive. The small rectangular footprint of the Sun StorageTek SL8500 modular tape system created by the U-shaped robotic design (Figure 5) offers very high floor-space utilization compared to traditional digital tape libraries. With a density of 50 tape cartridge slots per square foot, each library module is only 37.5 inches long, but adds another 1728 slots to the system. And the *Any Cartridge Any Slot* technology enables media production companies to consolidate multiple heterogeneous libraries and seamlessly mix tape drive technologies as they change and improve. As a result, there is never a requirement to replace slots or to add special drive frames to accommodate different technologies. And by consolidating multiple libraries into one Sun StorageTek SL8500 modular tape system, the library can also save on power and cooling. Similarly, the Sun StorageTek SL3000 saves floor-space, power, and cooling relative to other tape library products.



IBM TS3500 w/ LTO
19,797 slots; 180 LTO drives
 17 slots/ft² (185 slots/m²)
Note: library would be larger if included redundant robotics



STK L5500 w/ LTO
20,984 slots; 200 LTO drives
 25 slots/ft² (270 slots/m²)



Sun SL8500
19,500 slots; 192 drives
 50 slots/ft² (538 slots/m²)

Figure 5. The efficient design of the Sun StorageTek SL8500 library system results in optimal floor space utilization

Chapter 6

Conclusion

Long-term storage of feature films and broadcast programming has been a costly and frequently ineffective process. Rather than shooting analog film or video tape and converting it to digital for post-production work, the entertainment industry is embracing all-digital processes from the initial capture through post-production, distribution, and even projection. The convergence of these two dynamics, in addition to the economic and productivity benefits, is driving the media and entertainment industry to implement digital media content archives. Libraries of film and video content no longer need to be stored off-line in a large vault. They can be maintained in a single, scalable, digital archive system, helping production houses to complete projects on time and under budget.

Sun StorageTek hardware and software is designed to provide maximum scalability, performance, and throughput for the most content-intensive, media production environments. State-of-the-art tape libraries like the Sun StorageTek SL8500 and SL3000 modular tape systems and drives like the Sun StorageTek T10000 tape drive help make high-performance, high-capacity, and cost-effective digital archives a reality. In combination with Sun StorageTek QFS and SAM-FS storage management software, a digital archive can automatically manage the growth of most production houses' media content, eliminating the need for labor-intensive manual management of file systems as the content grows. In addition, user-defined policies determine the relative importance of content, and the system manages the data appropriately.

Among the numerous benefits of implementing a digital archive are the following:

- Enables longer-term storage helping to maintain content for decades
- Provides very high data integrity helping eliminate loss of content
- Supports virtually limitless capacity even for the largest production houses
- Enables faster and easier re-use of content helping complete projects on time
- Helps enable greater reliability and higher availability than digitizing and storing traditional content
- Offers a more cost-effective alternative to maintaining large library vaults

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Web site at: www.sun.com. Please also refer to the following resources for more information on topics discussed in this paper:

Web Sites

Information	Location
Sun StorageTek disks, tape drives, and tape libraries	sun.com/storage
Sun StorageTek SAM-FS and QFS software	sun.com/storage/software
Sun servers	sun.com/servers
Film degradation and restoration	en.wikipedia.org/wiki/Film_preservation
Process of digitizing a motion picture	en.wikipedia.org/wiki/Digital_intermediate
Film post-production processes and technologies	en.wikipedia.org/wiki/Post-production
Digital video recording systems	en.wikipedia.org/wiki/Digital_video
Library of Congress: Care, Handling and Storage of Motion Picture Film	www.loc.gov/preserv/care/film.html

Papers

Document	Location
<i>Sun StorageTek QFS and SAM-FS Software, Technical Overview</i> , March 2004	sun.com/storagetek/white-papers/qfs-samfs.pdf
<i>Archiving Movies in a Digital World</i> , White Paper, Dave Cavena et al, January 10, 2007	sun.com/storagetek/disk_systems/enterprise/5800/ArchivingMoviesinaDigitalWorld.pdf
<i>Sun StorageTek T10000 Tape Drive, Reliability Matters</i> , White Paper, May, 2007	sun.com/storagetek/tape_storage/tape_drives/t10000/488126_T10Kdrive_WP_LF.pdf

Organizations can also contact a local Sun sales representative to learn how Sun can help build competitive advantage with Sun StorageTek products that match data delivery needs from the workgroup to the data center.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN (9786) Web sun.com



© 2008 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, Sun Fire, StorageTek, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc. AMD Opteron and the AMD Opteron logo are a trademarks or registered trademarks of Advanced Micro Devices, Inc. Intel(R) Xeon(R) is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. Information subject to change without notice. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. Printed in USA SunWIN# 527097 Lit# STWP14023-0 02/08