An Oracle White Paper
July 2010

Digital Archiving Solutions for the Media and Entertainment Industry

ORACLE

Introduction

Almost since its inception, the motion picture industry has struggled to find an efficient and effective means of storing and preserving its most precious assets—the myriad feature films cranked out by Hollywood studios. Not only do both celluloid and videotape deteriorate over time, but both present problems of storage. For example, any videotape or film stored in a physical library or vault is far from readily accessible, often taking months to reach the requestor and entailing thousands of dollars in shipping and restoration costs. And if the content is to be reused in any new project, it will first need to be scanned and digitized. Finally, any storage that involves manual handling of film or videotape or variable environmental conditions increases the risk of damage to that content.

Increasingly, the media and entertainment industry is looking to digital content archives to solve these storage problems. With the industry embracing an all-digital process that encompasses everything from filming to postproduction and distribution, the time is right to move content libraries from offline vaults to scalable digital archive systems. This white paper helps organizations make that move by covering each of the four areas that must be considered in developing a digital archive—capacity, performance, reliability, and cost—and by describing Oracle's Sun servers, automated tape libraries, disk storage, management software, and infrastructure, which can be used to create and maintain such archives.

Challenges of Long-Term Media Content Archiving

Since the early years of the movie industry, filmmakers have struggled to find adequate long-term storage for motion pictures. Thousands of silent films were made in the years prior to the introduction of sound, and between 80 and 90 percent of them have been lost forever—many because of the deterioration of the actual celluloid film. During the first half of the twentieth century, movie images 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, which means that over the years, their negatives and prints have simply disintegrated. In addition, many early celluloid films have been damaged by dirt, scratches, tears, color fade, color change, or film grain noise. Those films might also be missing scenes and sound that were censored or edited out.

Similar problems occur with long-term storage of video programming. Videotapes consist of an acetate, polyester, or Mylar base coated with an oxide. That oxide is covered with a lubricant to minimize friction as the tape passes over the read/write heads of video recording and playback machines. Over time, the lubricant can evaporate or undergo a chemical change that renders 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 (the glue that holds the oxide to the base), the tape can develop clear spots where no oxide remains—again resulting in the loss of recorded material. In addition, the substrate can become brittle, making it too delicate to pass through a video player without breaking apart.

The extensive—and expensive—process of accessing and reusing media also presents a 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's stored, brought up to ambient temperature and humidity, packaged, insured, and shipped to the requestor. 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 US$10,000.

In addition, much of today's programming is created, produced, and distributed in digital form. This means that if an archived celluloid film or analog video is to be reused in a new production, it must first be scanned and digitized—another time-consuming and expensive process—before it can be digitally restored and integrated into the new media format.

Finally, the mundane task of manually handling videotapes or films can be both time- and space-consuming, and includes the risk of losing or damaging the content. Archived material can be mislabeled and misfiled, making it virtually impossible to locate in a large vault. Videotape can be dropped or carried near a strong magnetic field, rendering it worthless. Celluloid film can suffer heat shock when moved from a hot delivery truck into an air-conditioned studio—again potentially incurring irreversible damage.

Benefits of Digital Archiving

Virtually all the problems described in the preceding section can be eliminated with a digital content archive. Because digital media is replacing analog videotapes and celluloid film in all phases of the imaging process, many producers believe that it's time for the industry to fully embrace the digital process. Digital content can be stored electronically on or near the production system, making the content readily available for editing or reuse with none of the time, cost, and risk inherent in locating the material in a vault and physically transporting it to a production facility. Digital content also makes unnecessary the time-consuming and expensive film-out process for converting celluloid film to a digital intermediate. Digital films and video programs that enable re-editing and reuse are readily accessible at reasonable cost. As a result, many media houses are considering upgrading their production facilities to include digital archives.

Digital archives help ensure that content stored today remains available for viewing and reuse well into the future. The cinema industry requires that films be stored and available for approximately 100 years, and the video industry is storing programs for 50 years or more. Through the prudent use of best archival practices—which include periodically upgrading hardware and software as well as making copies of digital content at appropriate intervals—digital media archives can be expected to securely store content for 100 years or more.

Overview of Digital Archiving Solutions from Oracle

Oracle offers the servers, automated tape libraries, disk storage, management software, and infrastructure needed to create and maintain a digital media archive. As shown in Figure 1, at the core of any digital media archive is a high-capacity automated tape library providing the needed long-term storage, integrated with a high-performance disk system delivering digitized content for online editing and postproduction. The backbone of the system is a Fibre Channel interconnect, providing the 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.

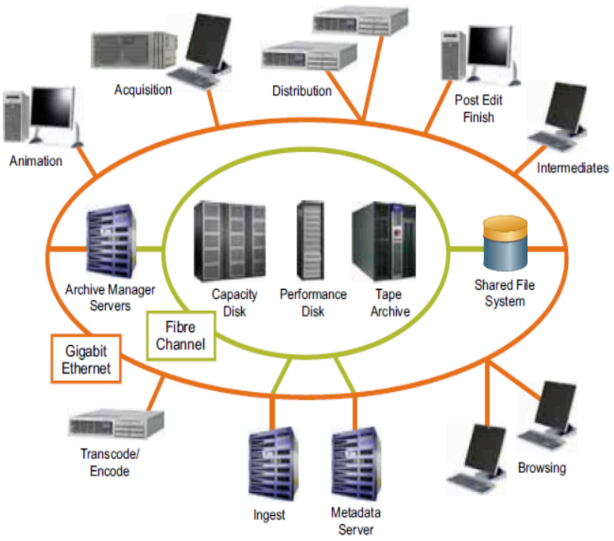

Figure 1. This illustrates a typical configuration for a midrange digital media content archive.

A typical archive system configuration includes

- **Long-term tape storage.** Oracle's StorageTek SL8500 modular library system, an extremely scalable and flexible tape library, is ideal for media archives. Using Oracle's StorageTek T10000B tape drives, the capacity of the library scales from 1.4 PB to 100 PB and supports from 64 to 640 drives. For environments that do not require these high levels of capacity and throughput, Oracle's StorageTek SL3000 modular library system scales from 200 TB to over 30 PB of tape storage; supports from one to 56 tape drives; and offers the same partitioning, advanced management, and redundancy features as the StorageTek SL8500 tape library.
- **High-capacity tape storage.** The StorageTek T10000B drive is a large-capacity (1 TB per cartridge), high-throughput (120 MB/sec) tape drive designed for extreme reliability in high duty cycle media production environments.

  High-performance first-tier storage. Oracle's Sun Storage 6780 array is a highly scalable and reliable disk array that provides exceptional performance. Designed to support up to 448 disk drives for a total of 268 TB of fibre channel storage capacity, the Sun Storage 6780 is highly adept at supporting throughput-intensive applications and its linear scalability ensures that performance scales to meet demanding applications and growth requirements.
- **Cost-effective production processing.** A wide range of highly scalable SPARC, Intel Xeon, and AMD Opteron processor-based servers can cost-effectively and reliably meet the performance demands of virtually any media production house.
- **Powerful storage management.** The combination of Sun Storage Archive Manager SAM and QFS software creates a tightly integrated data management solution designed to optimize a

tiered approach to storage hardware, providing innovative file system, file sharing, and archiving services.

This white paper focuses on the suite of components listed above because they offer 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 options for storing data, including optical disks, magnetic disks, CD-ROMs, and DVDs. Yet the computing industry's first reliable format for long-term data storage—digital tape—continues to play a major role in storage solutions 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 no matter what happens elsewhere.

Tape is the most affordable storage solution available. Because it is portable, tape can easily be stored in a safe and secure facility. Tape is dense—a single StorageTek T10000 tape cartridge can hold 1 TB of data—20 times the capacity of a standard Blu-ray optical disk. Tape is high performance: the StorageTek T10000 tape drive can transfer data at 120 MB/sec, more than 27 times Blu-ray's transfer rate. Most important, tape is safe, employing smart automation software and backup and recovery processes that are validated worldwide. Tape's legendary automation and virtualization capabilities have earned it the trust of some of the world's most successful companies. Today, with device-level encryption, a lost or stolen tape recorded on the StorageTek T10000B can be rendered useless to anyone who tries to read it without authorization (see Figure 2).

Figure 2. The StorageTek T10000B tape drive is a high-performance, cost-effective tape storage system.

Oracle 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. Oracle's tape libraries and drives are easy to install, scale, manage, and maintain. The 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 Oracle offer unprecedented durability and dependability. Highly accurate error correction code (ECC) helps ensure data integrity, while improved tape guiding systems designed to reduce tape speed with fewer tape passes help reduce drive and media wear. Oracle 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 in Assembling a Digital Archive

A digital archive is generally viewed as 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 white 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

Capacity

The volume of film and video content recorded by the media and entertainment industry each year is growing exponentially, with much of this content now being recorded digitally. All-digital motion picture and video production is becoming increasingly common due to the advent of digital camera technology suitable for shooting feature films. Capturing images directly in digital format simplifies postproduction work such as editing, sound, color correction, and special effects. And because digital content neither consumes film nor requires preprocessing to turn it into 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: The system must incorporate enough storage to accommodate all the digital content anticipated. Typically, an hour of high-definition 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, the required capacity of the archive can easily be estimated. However, other aspects of capacity must also be considered, including projected growth of the archive and accessibility of the content.

Storage space must be available to accommodate all productions scheduled for the foreseeable future. In general, however, it's unwise to project size requirements for an archive beyond two or three years, because the assumptions used for such projections are unlikely to be valid beyond that time frame. In addition, given the potential for changes in storage technology, what currently appears to be more than sufficient storage could be inadequate in five years. Finally, because unexpected and sometimes rapid content growth is almost inevitable, the archive should be scalable in easy and cost-effective increments.

Once storage requirements have been determined, storage type—online, near-line, or offline—should be considered.

- **Online storage.** Used for content that must be immediately accessible to production staff, online storage takes place on high-performance online disk arrays.
- **Near-line storage.** Used for content that needs to be accessible within a few minutes (rather than immediately), this type of storage can take place in a more cost-effective tape library. Keep in mind, however, that this library must be highly scalable because it must ultimately accommodate the entire archive, as well as any further unexpected additions.
- **Offline storage.** Used for content that doesn't need to be accessed in the near term, this type of storage takes place on tapes stored securely offsite. (Offline storage includes the physical space needed for the archive vault.)

Virtually Unlimited Near-Line Capacity

The StorageTek SL8500 modular library system is uniquely designed to provide virtually limitless storage capacity while allowing growth in easy-to-manage, cost-effective increments (see Figure 3). A basic StorageTek SL8500 library can accommodate up to 1,448 cartridges, providing 1.4 PB

[1] The DI process essentially converts a movie from celluloid film to digital bits for postproduction work, and then back to film for release to theaters. Although the size of a DI file depends on a film's length, for the purposes of this paper, a DI is assumed to be 10 TB.

(uncompressed) of near-line storage capacity and offering users access to as many as 144 feature film DIs.

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

Each 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. Because each additional library module contains its own four Handbot robotics, the system's performance does not degrade when capacity is added—as it does with many traditional tape libraries.

The StorageTek SL8500 library includes a minimum of 1,448 cartridge slots and can grow to 10,000 cartridges—for 10 PB or 1,000 feature films. Multiple 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 and providing up to 100 PB of storage—enough to provide rapid access to as many as 10,000 feature films.

Based on the innovative design of the StorageTek SL8500, the midrange StorageTek SL3000 modular library system supports up to 56 tape drives and offers nondisruptive scalability to over 3,000 slots—providing over 3 PB of storage capacity when used with StorageTek T10000B tape

drives and cartridges. Of course, with either library system, additional cartridges containing even more films can be stored offsite for later use.

Online Capacity for Production

The Sun Storage 6780 disk array stores up to 268 TB of digitized content (enough for 28 movie DIs) online and available for production work. The Sun Storage 6780 can be installed with just five disks and expanded as needed up to 448 disks, with a mix of high-performance solid state disk and/or fibre channel drives and high-capacity SATA drives in easy to manage, cost-effective increments. And, as previously discussed, all disk capacity upgrades can be conducted without disrupting the routine operation of the system.

Managing Massive Digital Media Archives

Once movies and videos are digitized and written into disk arrays and tape libraries, the task of managing and protecting these assets becomes paramount. Oracle 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, Oracle and its partners help media and entertainment companies easily and cost-effectively store, edit, reuse, distribute, and archive their digital media content.

Many media production houses are now investing in digital archive systems that can support multi-terabyte files contained in file systems of tens or even hundreds of petabytes. These capacity requirements strain the limits of traditional file systems and require a solution that has enterprise class scalability. The advanced file system that forms the foundation of Sun QFS and Sun Storage Archive Manager storage management software is designed to seamlessly scale as data storage needs grow, supporting virtually unlimited capacity. Furthermore, Sun SAM QFS 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. SAM QFS enables the user to shrink or grow without unmounting the file system and thus mitigating any associated downtime. 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 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 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 Storage Archive Manager (SAM) software is a hierarchical storage management (HSM) solution that reduce the cost to store content and helps improve the storage capacity use of the media production system, as well as the productivity of the production process, by enabling dynamic archiving and fast data retrieval. It automates the four core processes of Archive, Release, Stage and Recycle needed for a cost-effective and efficient archive architecture. 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. Also the archived data is never held "hostage" and mitigates vendor lock because all data is stored in a standard UNIX tar format that can be accessed, read and modified without the use of SAM QFS.

Performance

Traditionally, obtaining video or film content has required 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 take months. Today, digital archive users expect the system to deliver a stored program or movie to their desktops in minutes. The primary features that enable archive systems to do this are as follows:

- 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, which increases the effective throughput of the tape library
- Sophisticated file system software helping to optimize the handling of large data sets
- Online, high-performance disk storage to cache current production files for immediate access

Retrieve Content Faster

Although it is much faster than obtaining a celluloid film or videotape, retrieving a single stream of digitized content can still take a long time. Reading a 10 TB DI file on a StorageTek T10000

tape drive at 120 MB/sec takes approximately 23 hours, not including the time to load and unload the 20 or so tape cartridges involved. However, because the basic StorageTek SL8500 library supports as many as 64 StorageTek T10000B 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/hr, 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 Sun SAM QFS helping production facilities access their content faster and complete their projects on schedule. Because both the Oracle Solaris operating system 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 Storage Archive Manager 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.

Improve Archive Performance

With unique features for handling extremely large files such as video clips and digitized motion pictures, Sun SAM QFS software help enhance the performance of digital archives. Further improving read performance is the use of the MySQL Sideband database. The Side band data base offloads the indexing of metadata and can be storage on any storage device to deliver desired performance levels. 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 I/O patterns on the storage devices. These features help speed media content transfer between the tape library and the production system, helping the production staff complete projects sooner.

Sun 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 limits on the number of files in the file system. To increase throughput, Sun QFS software minimizes the movement of the physical head on the disk with new technologies such as metadata separation, variable block size, preallocation, and automatic direct I/O.[2] In addition, 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 archive's overall productivity.

[2] Metadata provides information to the file system about the characteristics of the data stored.

Sun SAM software consolidates innovative archiving and backup methodologies into its high-performance file system. The software replaces traditional backups to improve storage resource use 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.

Oracle Solaris helps Sun servers achieve superior performance and scalability by using threads as the fundamental technology employed to allocate processors.[3] Many threads can be executed in parallel by separate processors in multiprocessor systems. Oracle Solaris responds to interrupts, performs driver and background activities, and handles application requests using threads. The number of possible threads is limited only by the amount of available memory. In addition, SAM QFS takes advantage of Solaris Zones by virtualizing OS services and preventing the impact of concurrent processes occurring other zones. The net effect is reduced operating costs through easier management in virtual environments. By using the same mechanism for scheduling both system activities and user processes, improvements in threads enhance all aspects of performance for applications built to use multiple processors, including the advanced file system in Sun SAM QFS software. In an environment where both the operating system and file system support multithreading, multiple threads can access a file simultaneously—beneficial for applications such as media streaming and large media file production.

Increase Operational Performance

Due to the relatively lengthy operation required to obtain a digital file directly from a tape library, tape-only archive systems are not sufficient for production environments, 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 StorageTek T10000B tape drives operating in parallel. To resolve this, a high-performance disk storage system, such as the Sun Storage 6780 array, is coupled with the tape library so that digital content is read from the tape library and cached in the disk array, enabling postproduction functions to be performed in real time (see Figure 4).

[3] A *thread* is an independent sequence of program instructions that can be executed by a processor.

Figure 4. Designed to handle large data sets, the Sun Storage 6780 array delivers industry-leading performance.

Reliability and Availability

One of the major concerns of the media and entertainment industry with respect to digital archiving is ensuring that the production process can continue to operate in the event of a hardware or software disruption, and that the data captured in the production process remains unchanged once it is stored in an archive. The equipment used in a digital media archive must use state-of-the-art ECC technology to ensure precise data reproduction. In addition, storage drives and arrays must be designed with the necessary high-availability technology to ensure that content production continues even when a component fails.

Engineered for Data-Intensive Environments

The StorageTek T10000B tape drive and tape cartridge are designed to work in high duty cycle environments—operating 70 to 100 percent of the time.[4] Both the drive and the cartridge include advanced technologies that minimize wear, resulting in higher reliability. The StorageTek T10000B drive uses a unique, 32-channel, dual-head technology—enabling slower tape speeds,

[4] A d*uty cycle* is the amount of time that the device is in operation.

fewer passes to fill a cartridge, and better use of the tape media—which helps ensure reliable reading and writing. Slower tape speed—3.74 m/sec versus 6.2 m/sec 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 StorageTek T10000B tape drive uses a sophisticated dual ECC scheme and a cyclic redundancy check 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 reconstruct the data from the ECC bits recorded on the tape. The ECC scheme for the StorageTek T10000B 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 competing 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 dropouts and tape defects.

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

Sun SAM software includes a system error facility reporting mechanism, which helps notify system administrators if tapes have aged and should be replaced or rotated. In addition, Sun Storage Archive Manager 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 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 StorageTek T10000 cartridge has a unique locking mechanism that prevents the hub from hitting the inside of the cartridge and damaging the tape during rough handling.

Tape Drive Reliability

The tape guiding system of the StorageTek T10000B tape drive, further protects the media and the data. A 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 StorageTek T10000 tape drive uses a 32-channel, dual-head design that 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 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 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, allowing the library to continue operating while a failed power supply is replaced.

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

The real time growth capability designed into both the StorageTek SL8500 and the 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 increased 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 no need to shut down the system to replace slots to accommodate a different mix of media. Both the StorageTek SL8500 and StorageTek SL3000 library systems allow new drive and media units to be added, without installing unique slots to accommodate the different technologies. What's more, these performance and capacity upgrades can be performed while the system is operating.

Cost of Ownership

When evaluating whether to deploy a digital media archive or maintain a film or videotape 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 postproduction in using mostly digital technology, the cost of film and film processing will dramatically increase—if these alternatives 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. In addition, the cost and availability of laboratories to perform 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 videotape technology usually requires rerecording the content onto new, compatible tapes, with a resulting deterioration of video quality, because the copy retains 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 demise 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 floor space, power, and cooling. Archived content should be rerecorded to new tape media periodically to help confirm that the tape cartridges and drives are compatible. This not only helps prevent problems caused by media deterioration, but it also provides the opportunity to transcode assets to current, state-of-the-industry media and formats, which in turn helps ensure that the assets can be retrieved with currently available systems. Also, periodically rewriting the content keeps 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 5 to 10 years. Tape drives should be replaced in 3- to 5-year intervals and tape libraries in 10- to 20-year intervals. The published media life for the StorageTek T10000 tape cartridge is 30 years, but tapes should be replaced every 10 to 15 years to help ensure their 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 SAM software is a HSM software that determines where content is stored based on user-defined policies. Media content is distributed across multiple levels of storage such as Solid State Disk (SSD), 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 less urgently needed content on the more economical, near-line tape library, the 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 (and more costly) online storage devices.

The StorageTek T10000B tape drive and StorageTek SL8500 tape library are designed for high duty cycle environments—in many cases, operating more than 20 hours per day, 365 days a year. Not only does the system's high duty cycle design help ensure reliability, but it also helps 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 x 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 Oracle help reduce costs by eliminating the need to purchase additional lower duty cycle tape drives that would be needed to support the throughput requirements of large media files.

Controlling Media Costs

The StorageTek T10000 tape cartridges can be used for at least two generations of tape drives. Unlike typical competing tape drives that are read-only backward compatible to the previous generation, the StorageTek T10000B tape drives are designed to be to be fully compatible, saving operating expenses by enabling the cartridges to be reused with the next-generation drives. Because media often represents the largest expense in an archive, reusing it for a second generation can result in a huge cost savings.

To reduce the amount of media required, Sun Storage Archive Manager 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 StorageTek SL8500 modular tape system created by the U-shaped robotic design offers very high floor space use compared to traditional digital tape libraries. With a density of 50 tape cartridge slots per square foot, each library module is only 37.5 in. long, but adds another 1,728 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's never a need to replace slots or add special drive frames to accommodate different technologies. And by consolidating multiple libraries into one StorageTek SL8500 modular tape system, the library can also save on power and cooling. Similarly, the StorageTek SL3000 saves floor space, power, and cooling expenses relative to other tape library products.

Case Studies

The following case studies illustrate how two organizations are using Oracle hardware and software to create cost-effective, accessible digital archives for their digital media content.

Electrofilm USA GmbH

Electrofilm USA 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, and would enable managers to scale up capacity easily and control the entire data lifecycle with minimal human intervention.

Electrofilm chose a system that includes StorageTek 6140 and StorageTek 6540 disk arrays for primary and secondary disk storage, a StorageTek SL8500 tape library with StorageTek T10000 tape drives for archival storage, and Sun Storage Archive Manager and Sun QFS storage management software running on Oracle's Sun Fire X4600 server for content lifecycle management.

With these server and storage products, Electrofilm has created a postproduction system—based on a tapeless workflow—that uses no analog videotape 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 videotape, enabling it to compete more effectively with other companies that use a mix of analog and digital.

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 stored 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 to protect them from natural deterioration and to improve the public's access to these artifacts.

The center designed an archival infrastructure using Oracle servers and storage. The system includes high-performance Sun Fire X4600 servers running Oracle Solaris 10; a robust SAN network based on StorageTek 6540 arrays; and a StorageTek SL8500 tape library with StorageTek T10000 tape drives, managed with Sun SAM QFS 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 that stored content can be preserved. As a result, the public has easier and faster access to many long-forgotten audiovisual treasures.

Conclusion

Historically, long-term storage of feature films and broadcast programming has been a costly and frequently ineffective process. Today, rather than shooting analog film or videotape and converting it to digital for postproduction work, the media and entertainment industry is embracing all-digital processes from the initial capture through postproduction, distribution, and even projection. In addition to providing economic and productivity benefits, the convergence of these two dynamics 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 offline in a large vault. Today, they can be maintained in a single, scalable digital archive system, helping production houses complete projects on time and under budget.

Oracle's 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 such as Oracle's StorageTek SL8500 and Oracle's StorageTek SL3000 modular tape systems and drives such as Oracle's StorageTek T10000B tape drive help make high-performance, high-capacity, and cost-effective digital archives a reality. In combination with Sun QFS and Sun Storage Archive Manager 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, maintaining 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 reuse of content, making it easier for organizations to complete projects on time
- Delivers greater reliability and higher availability than digitizing and storing traditional content
- Offers a cost-effective alternative to maintaining large library vaults

Oracle is committed to developing practices and products that help protect the environment

Digital Archiving Solutions for the Media and Entertainment Industry
April 2010

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

SOFTWARE. HARDWARE. COMPLETE.