

Digital Archiving for Filmmakers

Best practices for safely backing up, storing and
managing your footage—from independent, nonprofit
and corporate to broadcast and feature production



studiodaily PRESENTS

Digital Archiving for Filmmakers

Editor
Beth Marchant

© 2010 Access Intelligence, LLC

studiodaily

www.studiodaily.com

Guidebook Editor

Beth Marchant
(212) 621-4645

Graphic Designer

Christina Barnard

Film & Video

EDITOR-IN-CHIEF
Bryant Frazer
(212) 621-4647

Studio/Monthly

EDITOR-IN-CHIEF
Beth Marchant
(212) 621-4645

BUSINESS OPERATIONS/SALES

VICE PRESIDENT/GROUP PUBLISHER
Diane Schwartz

DIRECTOR OF EVENTS LOGISTICS & MARKETING

Kate Schaeffer

SALES MANAGER

Todd K. Spitz
(727) 452-4282, Fax: (813) 207-5013

REPRINTS

Lyndsay Bahn
800-290-5460 or 717-399-1900

The Studio Group

110 William St., 11th Fl., New York, NY 10038
phone: (212) 621-4900 • fax: (212) 621-4635
For additional copies of this Studio Guidebook,
contact clientservices@accessintel.com,
888-707-5814

For additional copies, please contact
clientservices@accessintel.com or

800-777-5006



ACCESS INTELLIGENCE CORPORATE

4 Choke Cherry Rd., 2nd Fl.,
Rockville, MD 20850
phone: (301) 354-2000
fax: (301) 738-7581

CEO AND PRESIDENT — Donald A. Pazour

EXEC. VP, CFO — Ed Pinedo

EXEC. VP, HR/ADMINISTRATION — Macy L. Fecto

DIVISIONAL PRESIDENT — Heather Farley

VP, CORPORATE CONTROLLER — Gerald Stasko

VP, CORPORATE CIRCULATION — Sylvia Sierra

VP, PRODUCTION & MANUFACTURING —
Michael Kraus

SR. VP/CHIEF TECHNOLOGY OFFICER —
Robert Paciorek

Why Archive?

The digital revolution has brought many things to filmmakers and content creators, among them better, faster workflows, higher-resolution images and sound, and access, via file-based formats, to infinitely more data at any given time throughout the production and post-production process. And therein lies the problem: With so much content being created by so many, how it is stored, managed and saved for posterity is no longer a task limited to archivists in some dusty, climate-controlled warehouse. It is the responsibility of every filmmaker who creates an audiovisual record of what engages and entertains us, from news, historical events and the natural world, to popular culture and our collective imaginations. But for too many archiving is an afterthought or, in the very least, an acknowledged but far-too-confusing process their current budgets, storage infrastructures and project loads just can't address right now. That is, until it's too late. When incorrectly stored and protected originals are accidentally erased through human error, neglect, system failure or natural disaster, everyone loses out.

But consider this: Every year storage capacities increase and overall costs go down. As facilities get more storage for the money, doesn't it make sense to put that savings back into the archival process—whether to invest in the digitization of analog material or in more complex management software and systems and storage infrastructures to support ever-expanding repositories of online and offline content? Of course, as resolutions increase, from HD up to 4K and above, there goes the real estate. The rising tide of 3D production and post will only add to bloated file sizes. But it shouldn't be an either/or proposition. Archiving systems come in vastly scalable configurations that work in tandem with existing editing workflows and SAN, NAS and tiered storage. On the simplest level, digital archiving begins with naming and organizing your metadata, both in-camera and offline, for more efficient storage and retrieval. More complex digital archive and asset management systems have the potential to both save and make a facility money in the long term. What you archive today should be available indefinitely for financial, cultural or artistic gain.

This *StudioDaily Guidebook*, the second in our new series, outlines the components of digital archiving and how it can easily become a part of existing workflows for a range of content owners, from independent filmmakers and small production and post facilities on up to emerging broadcast facilities and film studios. We have collected the most topical case studies, articles, research, white papers and vendor resources into one report to help you connect the dots and better plan the best preservation path for your particular content. The object is to present the pros and cons of various archival media and, through specific examples, the variety of methods available to facilities with varying physical size, media stores and budgets.

Beth Marchant
Editor



The Archive Moves Center Stage

Digital content just keeps getting bigger. Cable operators are offering scores of channels in HD, high-profile filmmakers are shooting in 4K, and 3D and animation are not just for kids anymore. Entirely new delivery channels and platforms are emerging, from on-demand to online delivery to the iPhone.

The biggest impact is on the workflow: how to store and manage these vast collections of often very large files. With storage, it's always about balancing availability, performance and cost. But when your storage needs explode, can you be sure your digital content will be always available and instantly accessible—and that you can afford it?

At Atempo, more and more of our customers are turning to digital archiving to tackle these issues. Our customers are creating open, active archives that are accessed from all stages of their workflows—from capture and ingest, through edit, and into distribution. One of our customers is archiving 2.7 terabytes of dailies every hour, as soon as the footage enters their facility. Another customer gave editors direct archive access so they can easily repurpose content—and expanded their programming from 4-5 hours to 8-10 hours per week. A third is digitally preserving 70 years of film history, and is archiving 2.4 million files and 10 million metadata entries. And that's just for the first 100 titles.

These numbers may sound incredible, but to us at Atempo they're not surprising. Starting in 1993, we've delivered data protection and archiving solutions for enterprise customers worldwide, from Airbus to EFILM to ReelzChannel. Our technology roots run deep with some of the biggest names in storage, and we continue our tradition of innovation, building load-balancing, failover, deduplication and replication into our products.


Our commitment to media and entertainment means our software solutions are compatible with industry formats like MXF and DPX, and are integrated with workflow tools like Final Cut Server, Building4Media and CatDV. Our experience gives our products and our company the stability to be there for your organization and your content, now and for many years to come.

We are proud to sponsor this *StudioDaily Guidebook, Digital Archiving for Filmmakers*. We hope it provides you solid guidance for containing costs and delivering more revenue through your digital content, including strategies you can capitalize on today.



A handwritten signature in black ink that reads "Neal Ater". The signature is written in a cursive, flowing style.

Neal Ater
CEO Atempo



studiodaily
home of the studiodaily.com online family

studio^{/monthly}
shoot • post • deliver.

Covering Technology's Impact on the Creative Process
FILM & VIDEO

The**Plug-In**Post

Studio**Bytes**

studio/direct

studiodaily fix

studio's
filmfest
compete.watch.share.

studiodaily.com

CONTACT INFORMATION
Todd Spitz, Sales Director
tspitz@accessintel.com or 727-452-4282

Table of Contents

Chapter 1: Archiving 101	6
An Archiving Primer	8
An overview of relevant terms, trends and emerging archival media and systems <i>by Beth Marchant</i>	
Digital Preservation Strategies.....	12
Understanding the risks and threats to digital data <i>by The Academy of Motion Picture Arts and Sciences' Science and Technology Council</i>	
Archiving in the Post-Production Market	18
A survey of archiving and storage options, benefits and challenges within the film and television post-production community <i>by Thomas M. Coughlin</i>	
Chapter 2: Archiving in Action	26
The DIY Archive for Independents	28
Create a secure, low-budget archive using LTO-4 tape on the Mac <i>by Helmut Kobler</i>	
Digitizing a Museum Tape Archive	31
How the nonprofit HistoryMakers transformed 2,000 oral histories into a digital archive for generations to come <i>by Steve Kwartek</i>	
Archiving on Set: Production	35
Getting <i>The Group</i> together; plus, tips for better archiving on set <i>by Beth Marchant</i>	
Archiving on Set: Post	38
Authentic Entertainment's archiving workflow for <i>Naughty Kitchen</i> and <i>Ace of Cakes</i> <i>by Jessica Clegg</i>	
Creating a Shared Tiered Storage Archive.....	40
How RLTV kicked the videotape habit <i>by Bryant Frazer</i>	
Restoring and Archiving Film for DVD, Web, Mobile and Beyond.....	43
Canada's Historic Visual Vaults Prepare for the Digital Future <i>by Beth Marchant</i>	
Chapter 3: New Directions	48
Ready for PrimeTime: The Archive's New Leading Role in Media Workflows	49
<i>by Janet LaFleur</i>	
Chapter 4: Resources	52
Directory of primary manufacturers and companies offering archiving-related products and services. Includes brief descriptions of products.	53

All of today's
technology products,
including storage
media, hardware and
software, have a finite
lifetime, and the time
required to migrate can
exceed the data's lifetime.

1

Archiving 101

An Archiving Primer

An Overview of Terms, Trends and Emerging Media and Systems

by **Beth Marchant**

What, exactly, is an archive? A cultural heritage to preserve for future generations? A valuable asset to be exploited for financial gain? A little of both? In an era of tightening budgets, the archive has become an increasingly valuable commodity. Archived content could fill out a broadcast schedule when production gets cut. It could supply critical segments of footage to the backstory in a developing news story. It could solve historical, critical or even scientific puzzles. It could be rediscovered by future generations in a variety of formats on multiple devices, decades, if not centuries, from now.

If masters, on hard drives, of your last project are sitting in a shoebox in your closet, it's time to consider some alternatives.

The biggest problem, especially for independent filmmakers, is budget. Investing in archival software, systems or devices, too often gets ignored by far sexier necessities, like a new piece of editing software or a hot new camera.

Add to that the tapeless conundrum. The more content—and subsequent studios and business units—in question, the greater the need to manage the archive throughout the production and post process. At the high end, the cost of storing 4K digital masters is more than 1,000% higher than the cost of storing film masters.

But one thing filmmakers at every level can agree on: If you are the content owner, whether a small documentary studio or a large broadcast facility, you need to hang on to your footage in its original form. Compressed files may work in the short term, but they won't give you as many options in the long term. What approach you take depends on the size of your archive and, of course, your budget.

Tape in a Tapeless World

Front Porch Digital's CEO Mike Knaisch estimates that there are some six billion hours of videotape sitting on shelves just waiting to be digitized. As the tools evolve, expect to see even more labs and facilities known for other common post-production services to begin offering digitizing as just one more option on the menu. In May 2010, Technicolor began offering just such a service, aimed at the Hollywood film and television market, after it installed a Front Porch Digital SAMMA robot system (see Resources, Chapter 4, page 50).

Says StudioDaily contributor Helmut Kolber, "The problem with tapeless is that it requires rethinking how footage is shot and moved through post-production (and beyond). After 30-40 years of shooting tape, many people are reluctant to let go and learn something new (and suffer some of the slings and arrows that come from venturing into new territory...and there are definitely a few slings and arrows!). But the move to tapeless is in full swing, and it's only going to accelerate."

The not-so-surprising twist in a tapeless world? Whether a tiny, one-person shop, a museum or a much larger post or broadcast facility, the archival medium of choice is often LTO tape.

Linear Tape Open is a family of open magnetic tape standards developed by HP, IBM and Quantum that are licensed to third-party vendors. The half-inch magnetic tape has been used for data storage for more than a half century, with little sign of slowing down. DEC's DLT was an early variety of LTO. Even the ethereal sounding cloud storage—developed to archive e-mail and other non-critical data for the record but recently expanding its use to larger media files—uses a combination of disk and magnetic tape for its deep archive.

How Long Will It Last?

Most manufacturers of LTO-4 drives and tape, which were first introduced in 2007, and LTO-5, introduced in March 2010, confirm a 20-to-30 year shelf life. But perhaps equally important, each LTO format has seamlessly transitioned from one version to the next with backward compatibility built in. LTO Ultrium, the single-reel standard found in most LTO-4 and most recently, LTO-5 drives on the market, was designed to be a drop-in replacement for similarly sized, half-inch DLT tapes. Those with a substantial DLT archive could easily convert their tape libraries to the longer-lasting LTO format. The Ultrium can also read LTO tapes at least two generations older than itself and write to one from the preceding generation. Its developers say it is both “scalable and durable,” with a very long future ahead.

Why not just archive your masters on the tapes they were recorded on? The latest LTO-5 cassette will hold approximately 1.6 TB of data; LTO-4, about 800 GB. As contributor Helmut Kobler points out, “You’d need about 50 DVCPRO HD tapes to store the footage that a single \$40 - \$100 LTO-4 tape handles.” It’s not exactly practical to think about the kind of physical space you’d need to store hundreds of thousands of video tapes shot over the years, let alone how much it would cost to rent it. Solid state flash drives, which one day may hold vast repositories of content, are still too small at present to be considered archival material.

Although hard drives and disks continue to make impressive capacity increases—and become cheaper per TB as a result—you should never think of your magnetic hard drives as archival repositories, adds archiving expert Milton Shefter, the president of Miljoy, Inc., in Los Angeles. “They are designed to be powered on and spinning, not to be stored on a shelf for extended periods. They lock up if they aren’t in use.” Newer systems that combine magnetic tape and hard disk, such as Massive Array of Idle Disks (MAID), where disks spin down when not being used, are one solution. Copan Systems, acquired by SGL in February 2010, is currently co-developing MAID systems that work with SGL’s FlashNet open archiving system aimed at the broadcast market.

Holographic Storage

The digital devil on your shoulder is whispering, “This is tape, after all. How long term—even with endless capacity and transfer speed upgrades—can

it truly be?” The future-proof sounding Holographic Disk Storage (HDS) is another emerging archival technology, albeit one that has lumbered on the horizon for the past five years or more. At press time, HDS promises upwards of 1.6 TB on a single disk and promises a better price per TB, higher capacities, faster restore speeds and an expected shelf life of about 50 years. With medium-sized facilities increasing their storage, on average, by about 10-to-50 TB this year (large broadcasters and studios handling exponential daily content or more complex, higher-resolution files will likely increase by 200 or more TB), the technology is intriguing. Conceptually, holographic also makes more sense than tape: instead of recording data one bit at a time, holography records

The not-so-surprising twist in a tapeless world? Whether a tiny, one-person shop, a museum or a much larger post facility, the archival medium of choice is LTO magnetic tape.

and reads over a million bits of data, in its full three-dimensional depth, with a single flash of light. The benefits of this random-access technology include affordability, especially when compared to RAID, and near-line access to hundreds of gigabytes of data.

The most appealing promise of holographic storage for archiving? You will only have to migrate to a new format every 20 years. But the path to product for many holographic storage developers has been a rocky one. InPhase, one of the technology’s leading developers and one of the few to target the broadcast and film industry beyond the wider data enterprise market, only began shipping its first professional product, the Tapestry 300r drive and companion disk, in May 2009. The drive, with a 300 GB capacity and a 160 Mb/s transfer rate, lists for \$18,000; disks cost about \$180. InPhase had hoped to reach a 1.6 TB capacity benchmark with its first commercial drive in

2009. In fact, when InPhase first began to discuss its product strategy at trade shows and at industry events in 2005, the company expressed as much in its development roadmap. Unfortunately production hurdles, and the poor economy, intervened. The 1.6 TB goal could still be years away: The company now says it will release one more generation of product after the Tapestry 300r before it will release a commercial version with a full 1.6 TB capacity per disk. Holographic Storage, like LTO, is both WORM (Write Once, Read Many) and Encryption capable.

Is this enough to convince filmmakers to make one blind-faith leap—hands tethered and eyes closed—to Holographic technology? Not so fast, say experts at the LTO Consortium and comparable groups in the optical arena. The current spec for LTO-5 is already at 1.6 TB and has an uncompressed data transfer rate of 180 MB/s. Quantum, the first to release LTO-5 drives in early March 2010, claims that its drives can support up to 3 TB of capacity with transfer speeds up to 280 MB/s (based on a 2:1 compression rate). At the time of the release, Quantum also unveiled some of the newer features in LTO-5 technology, which include the kind of file control and data indexing that Holographic Storage offers. For example, you could set up a nearline application to index your LTO tape archive, making it easier to find your files. The backwards compatibility of the Ultrium format also means that this year's LTO drive and disk will be easily read by future LTO drives when they arrive.

LTO-6, on target for a 2012 release, is expected to support up to 3.2 TB natively, which will continue to shrink the cost of storage per GB. It's this economical advantage that has filmmakers—as well as storage and archiving manufacturers—relying so consistently on LTO tape to store and back up their media. In May 2010, magnetic tape reached an even more impressive milestone: Hitachi Maxell and the Tokyo Institute of Technology unveiled a magnetic tape cartridge that could hold 50 TB on a single cartridge.

These gains aside, The Academy of Motion Picture Arts and Sciences' Science and Technology Council points out that most leading tape vendors, as well as the National Archives and Record Administration (NARA), still recommend that digital assets on magnetic tape be migrated as “frequently as every five to 10 years.”

More Optical Options

In the motion picture industry, DVD-Rs are used more often than rewriteable Magneto-Optical (MO)

discs because they are less expensive and can store more per disc. But according to archiving expert Milton Shefter, Blu-ray Discs still “have a much smaller capacity per unit than LTO-4 magnetic tape.” “Gold Master” optical discs, which have a layer of gold in their recording layer, last longer but are five times more expensive than non-gold discs, he says.

“Optical storage technology hasn't kept up with magnetic storage technology in terms of areal density, capacity per unit, or transfer rates,” he says. There has been some progress in recent months, however. The Blu-ray Disc Association is releasing a new BDXL format that can hold up to 128GB (write-once) or 100GB (rewriteable). The downside? Because there are added layers on this denser disc, you'll need a new player, not just a firmware update, to play them back. If you are currently using optical discs for archiving the extra space will let you store much more per disc.

There have been some surprising recent product announcements in the optical disc arena. Primera and the optical disc company Millenniata, Inc., for example, made a startling announcement at NAB 2010 with the introduction of a new line of automated disc publishers. The claim? The DVD-R-type M-DISC you archive to will last, with your media intact, for up to 1,000 years.

If you run a solo shop or very small facility, this is a remarkable option for the price. An M-DISC enabled Bravo SE from Primera lists for \$2,995. Primera's BravoPro Xi2 lists for \$295. But one M-DISC, available now for about \$23, can hold only 4.7 GB, a fraction of what can fit onto one LTO-5 tape. And while the disk may outlive you and your next seven or eight generations of descendants, what about the players? No such claims can possibly be made about machines with any mechanical moving parts. If the optical disc format eventually disappears, finding a player to view what's on that M-DISC in 800 years could definitely be an issue.

“I find that claim impossible to believe,” says Front Porch Digital's Mike Knaisch. Yet, he adds, “the sense of urgency in any migration effort is so keen. We know we have a huge problem and mounting content, but we also know that over time, the playback devices themselves are perishable. Migrating, whether from optical disc to optical disc, or to a newer LTO format, is the only viable option right now.”

As we move into an era of more open formats, Milton Shefter, for one, believes that filmmakers and content owners need to work closely with their

archivists to bring the potential of the archive to fruition. “You can’t just hope to keep temperatures down and humidity levels low wherever you’re storing your magnetic tapes,” he says. “You’ve got to actively manage them,” something you can now do globally with a file-based workflows. “But that also means that the people involved, from producers to users to archivists, have to work together to get the most out of that archived content.”

Without a company-wide archiving plan, he and many others warn that “islands of archiving” will pop up among decentralized facilities or different business units—for example, in marketing and production, or among satellite studios—that create and archive their content using multiple formats and methods. That, all agree, creates time, money and energy-wasting redundancy.

Storage and the Digital Asset Manager

The archive is just one part of a complex managed pipeline that includes storage, asset managers and access to content and a variety of levels. Cloud storage is gaining momentum. Another increasingly popular system is tiered storage, which lets companies reduce total storage costs by assigning different content, or data sets, to various storage mediums and devices. More detail on storage in post production in particular is discussed in the third article in this chapter, on page 17. To read about tiered storage in use in television production, turn to page 39.

Larger media organizations often turn to some form of Hierarchical Storage Management, or HSM, to manage file backup and archiving over time. Whether sent to RAID, optical disc, or tape, each process has different costs and retrieval speeds that can be managed efficiently. As media ages, HSM systems often move files to deeper archives on slower but also less expensive types of storage. The problem with HSM systems, however, is they are often made up of proprietary hardware that only works with specific file formats. Software-based, open-architecture storage and archival systems from companies such as Quantum, Atempo and others are not bound by those same restrictions.

There are a variety of interchangeable terms used to describe the control of media files across your pipeline, including Digital Asset Management, Media Asset Management and Content Storage Management. On the most basic level, these managers typically ingest

In May 2010, Hitachi Maxell and the Toyko Institute of Technology unveiled magnetic tape cartridges capable of holding 50 TB of data.

and log, index and catalog, and search and retrieve. Robust editing functions, including a host of collaborative and review and approval features, typify more recent asset management systems on the market. Not all asset management tools ever did or do supply every piece of the puzzle. In the past, technology partners (and often, competitors) would come together to offer an end-to-end system for customers. Recently, open standards and metadata-centric workflows have given rise to new product categories from companies known for their cameras and editing systems. At NAB 2010 both Sony and Avid showed similar service-oriented architecture, or SOA-based, open media management systems (Sony Media Backbone and Avid Integrated Media Enterprise) that, through one interface, connect their products and third-party software and hardware for seamless control of footage throughout its lifecycle, from ingest to the archive.

The following articles in this chapter discuss the risks and threats to storing digital data and strategies—across the production and storage pipeline—for archiving and preservation. ■

Beth Marchant is the founding editor-in-chief of Studio/monthly and the co-editor of StudioDaily.com, a collection of Web sites, including Studiomonthly.com and filmandvideomagazine.com, dedicated to production and post-production technology. She has spoken on a range of digital content creation topics at the National Association of Broadcasters convention, the Film Finance Forum and at the Writers Guild of America. Her articles and photographs have appeared in Yankee, Harvard Magazine, on the syndicated Gannett News Service and in numerous Studio Group publications in print and online. She also reviews consumer technology for the site Techlicious.

Digital Preservation Strategies

Understanding the Risks and Threats to Digital Data

By the Academy of Motion Picture Arts and Sciences' Science and Technology Council

Two questions are key to understanding why digital archives cannot be preserved over the long term using a “store and ignore” management philosophy: “Is there any way to store a digital object for 100 years with no maintenance?” Secondly, “Is the bit density enough to hold what you want to preserve at a price you can afford?”

If one could make a “black box” with even 100-year lifespan components that could read data reliably without introducing any errors, required no maintenance, and offered sufficient bit density at an affordable price, everyone would buy it. After filling the black box with their most valuable “permanent” assets, one of the first things prudent archivists would do is create several replicas in multiple black boxes and geographically separate them to guarantee viability and enable the archive to become self-healing. If the archive format preserved both bits and needed application software together with contextual metadata, there would be no need for periodic data migration or system emulation. But there’s a new danger inherent in this approach. If the 100-year-lifespan “box” fails at 99 years, no one involved in its development or capable of system repair is likely to remain alive. To avoid this risk, it would be necessary to continuously audit the integrity of the box to ensure that the archived assets can move to a new box before the old box fails. This points to the need to sustain a supportive human community around a digital archive with the requisite know-how in order to ensure its ability to preserve, renew and repair the system within which digital assets are stored.

Digital assets in the real world are not kept in “black boxes” with 100-year longevity. They are stored on physical media with longevities of 30 years or less, and are vulnerable to heat, humidity, static electricity and electromagnetic fields. The

digital contents can be degraded by accumulating unnoticed statistically occurring “natural” errors, by corruption induced by processing or communication errors, or by malicious viruses or human action. Digital media cannot be viewed with the naked eye. As such, it is susceptible to misidentification, frequently poorly described (incomplete labeling and metadata), and therefore difficult to track. And digital assets are hard to maintain long-term because media, hardware and software can all become obsolete. This is commonly caused by the evolutionary loss of compatibility between data in the archive and the software applications that originally created the data. Sometimes proprietary formats in an archive are simply abandoned when a company goes out of business. A digital archive may have many “layers,” each with its own finite lifespan. When the end of the lifespan is reached, not only does the layer have to be replaced, but the adjacent layers may have to be modified to be compatible with the replacement layer. Thus, a digital archive built with today’s digital technologies can only assure digital “permanence” via an ongoing and systematic preservation process. The rapid and seemingly endless improvement in the price per bit of digital data storage tends to give the impression that storage is forever getting cheaper, so why worry about the “data explosion”? There are several reasons why the overall storage picture is not as simple as this might make it seem.

Increasing Demand for Storage Offsets Reduced Media Cost

Along with the increase in available storage comes a corresponding increase in the demand for storage. In the UC Berkeley digital data generation study discussed earlier, it was found that of the 5 exabytes of new data created in 2002, 92% was recorded on

magnetic media, 7% on film, and the remaining 1% split between paper and optical media. Overall, UCB researchers estimated that new stored information grew about 30% from 1999 to 2002.

From the relatively narrow view of the motion picture industry, one only need consider the amount of data generated by 4K digital motion picture cameras and digital post-production process (in the petabyte range) to understand that there will always be a way to generate more data, usually in excess of available storage. The demand is compounded by the need to duplicate important data for backup purposes.

Data Transfer Rates Do Not Increase At the Same Rate as Storage Density

As the storage density grows, the speed at which the data gets on and off the storage media (transfer rate, or throughput) becomes more important. The need for increased throughput drives up the cost of the physical interface, network connections and computers attached to the disk drives. As with the demand for increased storage, throughput requirements increase with the need to make backup copies of important data.

Longevity Characteristics Aren't Always as Advertised

Studies by Google¹ and the Computer Science Department at Carnegie Mellon University² present evidence that hard drives are not as reliable as manufacturers' data sheets suggest, nor do they follow the conventionally accepted "bathtub curve"³ failure characteristic. To the contrary, these studies observe that large numbers of drives fail well before manufacturer-specified "mean time before failure" (MTBF), and show a low correlation between drive failure rates and high temperatures, a commonly assumed failure predictor.

The manager of a large digital image archive who has purchased a great deal of both tape and disk over the years said that in his experience, the biggest problem with a magnetic hard drive is its short device life cycle, supposedly five years according to manufacturers, but only three years in practice. He recognizes that disk technology is driven by personal computing and consumer electronics markets characterized by very short product life cycles, so there is naturally quite a bit

of product churn. In contrast, data tape drives are industrial products, with multi-year life cycles, and with some degree of backward compatibility and forward-looking roadmaps from vendors.

These empirical observations raise questions about the "accelerated age testing" methodologies used by storage product manufacturers to determine the life expectancy of their products, and suggest that there is no way of knowing whether a storage device or medium will, on average, last for the advertised period of time without actually seeing what happens during that entire time frame. It is worth repeating that both storage technology suppliers and end-users significantly de-rate the published life expectancies of all digital storage systems, usually planning on wholesale equipment and media replacement after as little as three years, with five to ten years as the most often quoted migration period.

Economic, Technical and Human Threats

A recent report by the National Research Council written for the National Archives⁴ presents the notion of threat modeling and threat countering as a core consideration in the design of digital preservation systems. These threats are further detailed in a paper on digital preservation system requirements published by the Stanford University Libraries⁵, and are worth summarizing here for the benefit of those responsible for preserving digital motion picture assets.

Economic Threats

Funding loss: Digital preservation systems require ongoing funding for equipment maintenance, replacement, operating staff and power, among other things. Every commercial enterprise has its good and less-than-good years, and the occasional "benign neglect" that film archives can tolerate may result in data loss in a digital archive. [Note: There is no known tactic to fully mitigate this threat, although factors that affect the economy of operating a digital storage system on a large scale for motion picture preservation are discussed in depth in the full "Digital Dilemma" report.]

Technical Threats

Data integrity: At the most basic level, the 0s and 1s that represent digital images and sound must be reliably stored and retrieved. Common failure modes that affect the integrity of the 0s and 1s pre-

The goal of archival data migration is preservation of the full information content, not just the bits.

served in digital archives are latent errors (errors lurking undetected), ingest errors (translation errors when digital data is brought into a digital system), and network communication errors (errors caused when digital data is moved between computers on a network). Regular auditing and authentication of the data and rigorous quality control procedures are effective means for dealing with these threats.

Monoculture vulnerabilities: Just as a single animal species can be wiped out from a deadly virus, individual storage media or technologies can be (and have been) seriously impacted in the same way⁶. Biodiversity, or the practice of utilizing several different media and technologies for digital storage, significantly reduces this threat^{7,8}.

Single point-of-failure: Storing a single copy of data in just one location is dangerous. Storage solutions should include sufficient redundancy to protect from data loss resulting from the failure of media, hardware, software, network services and/or natural disasters.

Obsolescence: All of today's technology products, including storage media, hardware and software, have a finite lifetime, and the time required to migrate can exceed the data's lifetime.

Limited or no data compression: A popular technique for reducing storage and transmission bandwidth needs is to apply mathematical data reduction techniques to image and sound data. These techniques range from "mathematically lossless" (every single bit is recovered when decompressed), to "perceptually lossless" (not every bit is recovered, but one cannot see or hear the difference between the decompressed content and the original), to "lossy" (perceptual artifacts exist in the decompressed content). The effects of compression must be well understood if used.

No risk of encryption key loss: There is much discussion today on safeguarding digital content through the use of data encryption methods. All encryption schemes require a digital key to "unlock" the encrypted content. If encryption is deemed necessary, then steps must be taken to eliminate the risk of losing the key, which is tantamount to losing the content it is intended to unlock. In general, there is broad consensus among those interviewed for this report that encrypting digital archives increases long-term complexity and risk.

Human Threats

Today's technology requires human involvement in many aspects of digital storage system operations. And being human means mistakes can and will be made. Furthermore, systems can be attacked by disgruntled employees or hackers simply doing it for fun. Procedures for protecting against losing media, unauthorized internal and external system access, reliance on a single employee, and storing multiple copies of important data in separate locations not controlled from a single place can be effective in managing the human element. Documentation of procedures and system implementation details can also protect against organizational failures that often occur when companies are sold or merged, or when key employees move on.

Broadly speaking, digital archiving experts have identified several preservation strategies that address either the general survivability of digital data or technical obsolescence. Two of those strategies are discussed here: migration and emulation.

Migration

Data migration involves the transfer of data from old physical media to new physical media, a process that often (but not always) includes updating file formats for currency with the latest-generation operating system and/or software applications. Older digital assets that are properly migrated will be accessible for some time into the future, until technological obsolescence motivates another migration cycle. Migration is designed to avoid having to preserve old devices to read the old storage media, old application software to interpret the old data, and old hardware to run the old software to use the old data. If everything goes smoothly, after migration the new data replaces the old data.

A major drawback to migration is that while copying data from one physical medium to another, or while converting digital assets from one file format to another, some data (or related metadata) might be lost. To make data migration a lossless, errorless process, migration procedures typically incorporate various quality control and auditing routines to ensure accuracy, integrity and completeness of the data throughout the migration process. Systemization of the migration process, including policy-driven automation routines, reportedly can be effective in reducing human errors and increasing the speed of migration. In practice, the emerging trend is to “migrate all the time” as a background task. Migration of archived assets by replicating them on new media is a preservation strategy for both analog and digital assets. An advantage of migration as a digital preservation strategy is that digital assets will always be available in the form that is most widely accepted, and current hardware and software will be able to render these digital assets with little difficulty. In the case of analog assets, migration can cause the loss of image and sound quality over successive generations. In the case of digital archiving, data migration done correctly is lossless every time. Data migration can occur between instances of the same type of storage medium, from one medium to another, and from one format to another. Data migration can be effective against media and hardware failures. For example, the tape backup of the contents of a magnetic hard drive involves data migration between different mediums.

The goal of archival data migration is preservation of the full information content, not just the bits. For example, the Open Archival Information System (OAIS), pioneered by NASA and others, defines “preservation description information” that should be included in the data migration process. This includes provenance information that describes the source of content, who has had custody of it, its history, how the content relates to other information outside the archive, and fixity information that protects the content from undocumented alteration.

Data migration can be motivated by a variety of factors such as physical media decay, media or media drive obsolescence, even prior to complete system obsolescence. Older media drives may face escalating maintenance costs, there may be new user service requirements, or new media formats and/or file formats are introduced that are more compatible with users’ technology and applications. The list of motivating factors goes on, and

therefore data migration is the most widely practiced digital preservation strategy today.

Emulation

Emulation preserves the original data format, often on the original physical medium, and provides the user with tools that enable the data to be read even after the original file format, storage medium, application program or host hardware is no longer supported. Emulation refers to the ability of one system or device to imitate another system or device. In practice, emulation involves writing software that runs on new hardware to make it appear as if it is an old system, translating between the two, allowing old data on old media to be “tricked” into working on a new system after the old underlying system has become obsolete. For example, new storage devices added to existing digital storage systems are often built with the ability to emulate an older storage device, so that the new storage technology can be integrated into the pre-existing software control and automation infrastructure of the system, thereby hiding the evolution of the infrastructure from the end-user. Emulation strategies for digital preservation are designed to minimize the need to copy, transfer, transform or otherwise “update” the digital assets in an archive. Digital archivists can use emulation strategies to reduce or even (theoretically) eliminate data migration. However, a serious drawback to emulation is the cost and complexity of developing and maintaining emulation tools. To avoid the risk that old emulation tools will not work on future computer platforms, software engineers must keep adapting and updating them.

While emulation has not been widely adopted as the primary digital preservation strategy for major digital archives to date, researchers at the University of Michigan and the University of Leeds in the UK, working with the BBC on the Domesday Project, have demonstrated that emulation can preserve the consumer’s experience of interactive multimedia based on older videodiscs and CD/DVD-ROM systems. They point to the need for emulation techniques in any effort to archive video games and hyperlinked rich-media documents.

This has led researchers, particularly some from IBM, to propose emulation strategies for long-term preservation based on the concept of a “Universal Virtual Computer” (UVC), a layer of software that remains the same on the “top side” facing the emulation

tools while evolving as needed on the “bottom side” facing the hardware and operating system (OS) software to adapt to changes in technology. In this approach, digital asset data is archived with a very basic software program that decodes the data and returns the asset in a readable form using a future software application based on a logical view that is simple and self-contained enough to be interpreted without any specific software or hardware. Working with the National Library of the Netherlands, IBM has successfully shown a proof-of-concept of the UVC approach using electronic documents deposited in the library in the Adobe Acrobat electronic document format⁹. Some argue that emulation, and its distant cousin encapsulation, are just more complicated forms of data migration.

No One Strategy Is “Best”

In considering emulation versus migration, experts agree that no one strategy is “best” for long-term preservation of digital data. Both emulation and migration have pros and cons. In general, storage vendors have tended to promote migration, while computer and software vendors have tended to promote emulation. Some digital preservation researchers advocate a hybrid approach, combining both migration and emulation. For example, emulation uses a “root format” from which digital asset transfers and conversions can be generated even as hardware and software evolve. But sometimes new formats are just too attractive to pass up, so an archive might periodically migrate its data to the new better/faster format, which then becomes the new root format for subsequent emulation. Among operators of major digital archives we interviewed, migration is the overwhelmingly preferred strategy for digital preservation at this time. But these same experts recognize that emulation also has merit, and admit emulation has been under-explored as a strategy for long-term preservation. Perhaps migration is the more conservative strategy and emulation requires higher initial investment in software development.

Sources:

1 Pinheiro, Eduardo, Wolf-Dietrich Weber and Luiz André Barroso. “Failure Trends in a Large Disk Drive Population.” Paper presented at the 5th USENIX Conference on File and Storage Technologies – Paper, February 12-13, 2007: 1-13 San Jose, California. 8 Aug. 2007 http://209.85.163.132/papers/disk_failures.pdf.

2 Schroeder, Bianca and Garth A. Gibson. “Disk Failures in the Real World: What does an MTTF of 1,000,000 Hours Mean to You?” Paper presented at the 5th USENIX Conference on File and Storage Technologies – Paper, February 12-13, 2007: San Jose, California. 26 Mar. 2007 http://www.usenix.org/events/fast07/tech/schroeder/schroeder_html/index.html.

3 The bathtub curve, used in reliability engineering, predicts early “infant mortality” failures, followed by a constant failure rate during a product’s useful life, followed by an increasing failure rate. A graphed curve of these failure characteristics looks like a bathtub – high at the ends and low in the middle.

4 National Research Council of the National Academies. Building an Electronic Records Archive at the National Archives and Records Administration: Recommendations for a Long-Term Strategy. Ed. Robert F. Sproull and Jon Eisenberg. Washington, D.C.: The National Academies Press, 2005.

5 Rosenthal, David S., et al. “Requirements for Digital Preservation Systems.” D-Lib Magazine 11, (11) 2005. 19 June 2007 <http://www.dlib.org/dlib/november05/rosenthal/11rosenthal.html>.

6 Herman, Robert. Telephone interview with Laurin Herr. 1 Feb. 2007. Oakdale, MN, USA.

7 Baker, Mary, et al. “A Fresh Look at the Reliability of Long-Term Digital Storage.” Proceedings of the 2006 Eurosys Conference, April 18-21, 2006: Leuven, Belgium. ACM Press, New York, NY, 221-234. 22 Jan. 2007 <http://www.lockss.org/locksswiki/files/3/30/Eurosys2006.pdf>.

8 Science Applications International Corporation. “Architecture and Technology Program: Offline Archive Media Trade Study.” United States Geological Survey. 2006: 1-34. USA. Aug. 2007 <http://wgiss.ceos.org/archive/archive.pdf/FY06%20Media%20Trade%20Study.pdf>.

9 Lorie, Raymond. “The UVC: A Method for Preserving Digital Documents: Proof of Concept.” IBM/KB Long-Term Preservation Study Report Series Number 4. International Business Machines Corporation and Koninklijke Bibliotheek. Dec. 2002: 1-47. Amsterdam, Netherlands. 18 Aug. 2007; http://www.kb.nl/hrd/dd/dd_onderzoek/reports/4-uvc.pdf.

This article was excerpted from “The Digital Dilemma,” used by permission and published by the Academy of Motion Picture Arts and Sciences’ Science and Technology Council.

STRATEGIES FOR INDEPENDENTS

Questions for Milton Shefter, Archiving Expert

Milton Shefter is a leading archival consultant to the motion picture industry, a former Hollywood SMPTE chair and president of Miljoy, Inc., a consulting and management company involved in media asset protection and preservation, including archive, library and warehouse operational design and environmental controls. He is best known for the creation and design of the extensive Paramount Pictures Asset Protection Program and his clients include Twentieth Century Fox, Universal Studios, Savoy Pictures and The Library of Congress. Beginning in 2007, Shefter was a lead researcher on the Academy of Motion Picture Arts and Sciences' Science and Technology Council's "The Digital Dilemma." Shefter will lead a follow-up report, to be published in late 2010, for the Council and The Library of Congress that is investigating the archiving options, practices and challenges for independent filmmakers and content creators outside the studio system. We asked him to share some of his findings thus far.

Q: What have you found out about the archiving habits of this group?

MS: *So many of these folks are only focused on getting to point of first exhibition. They are assuming, in some cases wrongly, that the distributor will take care of archiving their film beyond that point. However, many of them, documentarians in particular, do take their master material and put it in some form of cold storage, from closets to garages or someplace better equipped and climate controlled.*

Q: Are there any formal preservation programs out there right now targeted at this group?

MS: *Nada. Matter of fact, when we explain to them that there is an issue here, that their files may not be accessible to them for an infinite period of time, they are kind of shocked. Interestingly enough, one of the groups we've interviewed for this report is the independent distributors—who would not talk for publication, of course. We asked them what their contracts stated, regarding the elements that needed to be provided to them by the filmmaker. All of them said that their contracts had a specific clause that the content creator could ask to have whatever they had supplied in master elements back, anywhere from six weeks to six months before the end of the contract. Contracts usually run about five years. No one we spoke to could give me one example of a filmmaker asking for their material back. The filmmakers thought the distributors held onto it. The distributors, when asked, figured that the lab had the master, and so on, and so on. It's a finger-pointing issue. It's an issue we have to start facing, however, otherwise we're going to have a huge hole in our cultural visual history.*

Q: What's needed to solve this problem?

M.S.: *The archives are the end of the food chain. A lot of them are getting this digital material and have no idea what to do with it, don't know how to handle it, and certainly aren't funded. The issue as I see it, and what we're finding through our research, is that most of the vendors in this area do not believe the motion picture market is a big enough industry, in and of itself, to do anything for. The first thing that's needed is a cross-industry collaboration, put forth not only by the motion picture industry, but by government, big science, business. These are the guys that could make this thing happen. Number two: We must have some kind of standard. That's what missing. Every new standard has to be backward compatible. And we need an interim solution, and we need it now.*

—B.M.

Archiving in the Post-Production Market

A Survey of Archiving and Storage Options, Benefits and Challenges Within the Film and Television Post-Production Community

By **Thomas M. Coughlin**

Solutions and products for post-production must meet the needs of modern post-production facilities. The requirements differ depending upon the size of the organization. It is typical for a post-production house to have an active storage hierarchy that may be a combination of network storage, direct-attached storage and also a working archive for occasional backup of active projects. Due to the size of the content handled by these facilities there is usually not enough bandwidth in the network to support both real-time content access and hierarchical storage management.

Nonlinear editing (NLE) is generally done with uncompressed or at most slightly compressed content since heavy compression increases the overhead of editing and can cause timing problems. For a large enough facility with several editing chairs, shared storage allows the local disk storage to be kept at roughly 30 minutes per station. Use of storage networking exists for all the NLE market segments but is more likely for the higher-end market segments.

An archive is a copy of data that is being retained for very long periods of time, usually for years and in some cases, centuries. Archives are used throughout the media and entertainment industry, including post-production houses, for storing content that is not being used in ongoing projects, but could be re-purposed or referenced in the future. An archive can be active (online, where it can be accessed relatively quickly) or cold (offline, where it can be stored safely and economically, but may take time to mount the digital storage medium and read the archived content). The average time to access archived data is known as archive latency.

Editing and some other content industry segments also keep working archives of content on storage networks during the course of their work. These working archives are raw and edited content that are protected during active work on a project. The content in a post-production house is usually the property of the content creator rather than the post-production house. Archiving for post-production is usually done to provide data protection or to create near-line storage that contains data that still has value for the content creator and the post house, but does not require immediate real-time data access.

The working archives in post-production houses are often kept in storage area networks (SAN) or network attached storage (NAS) systems used in the working studio. After a project is completed, the content of a working archive may be retained in a long-term archive depending on the budget available, the value and time to create intermediaries, certain effects, and other content used to develop the final cut.

The file sizes in post-production can be very large since working content is usually raw content and therefore, it is difficult to perform traditional IT-style backups. Most production houses make periodic copies of projects to provide some level of disaster recovery.

Archiving for the Near and Long Term

Different facilities keep content for varying periods of time. Keeping completed content in long-term archives is common practice by content owners, including movie and television studios. Long-term raw content retention is not so common and

varies depending upon the policies and budgets of the facilities. Post, special effects and computer content generation houses may keep some of their unique content for extended periods. Long-term retention by post houses is not so common unless this is contracted by the content owner. As the cost of archiving declines and networking bandwidth increases, more content will be retained by all of these facilities.

The chart, above, shows projections for the costs of storage for 1 PB of content starting in mid 2008 over a 20 year period, assuming a hardware refresh (new HDD arrays or tape media) every 5 years and scaling other costs as appropriate with the decreasing number of HDDs (in a storage array) and tapes

small volumes of raw and edited content must make choices in what they preserve for the long term because they often do not have the resources or budget to save everything.

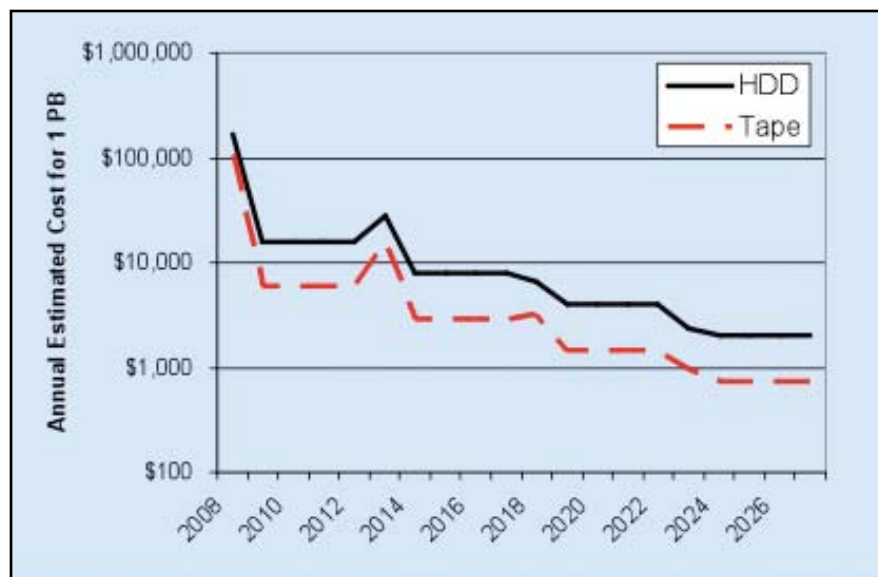
Preservation of source media containing the original content is almost always a requirement since it is difficult, if not impossible, to recreate this content. Preservation of the raw source material is often insufficient since raw source material does not capture any edits or metadata generated during the processing of the raw content to create a finished product. As a result, more content, in addition to the raw source material, may be archived.

To this end, many movie studios, with the help of their editing facilities, will save a set of masters (with color separation preserved) as well as managed digital tapes of the content in a cold archive. In addition, many of these facilities will also keep copies of content online in an active archive on a disk array or tape library for some period of time.

Smaller post-production shops and other facilities may not have the resources to set up a complex and expensive archiving system. For these users it is advantageous to retain the source media and players to play them back later. Careful logging and management of the physical media, combined with retention of the most

valuable content in disk arrays or tape libraries, may be a more affordable option.

Comparison of Estimated Annual Cost to Save 1 PB for 20 Years



Source: 2010 Digital Storage for Media and Entertainment Report, Coughlin Associates

in a tape library needed. The total cost of saving 1 PB of data for 20 years starting in 2008 is about \$320,000 for HDD storage and \$170,000 for tape storage. Over 50% of the total cost for both HDD and tape storage is spent in the first year. By the last year the annual costs of storing 1 PB is at most a few thousand dollars.

For preservation of new content, deciding what to archive can be difficult. The retention of digital post-production content involves cost in real estate, operation, management software and hardware. In practice, data managers handling large and even

The Role of Storage in the Working Archive

Post-production uses several types of storage with different performance and cost characteristics. Online storage must deliver real-time video using either high-performance direct-attached storage devices or high-performance network storage. This sort of storage has very fast data access and delivery but is also very expensive. Therefore,

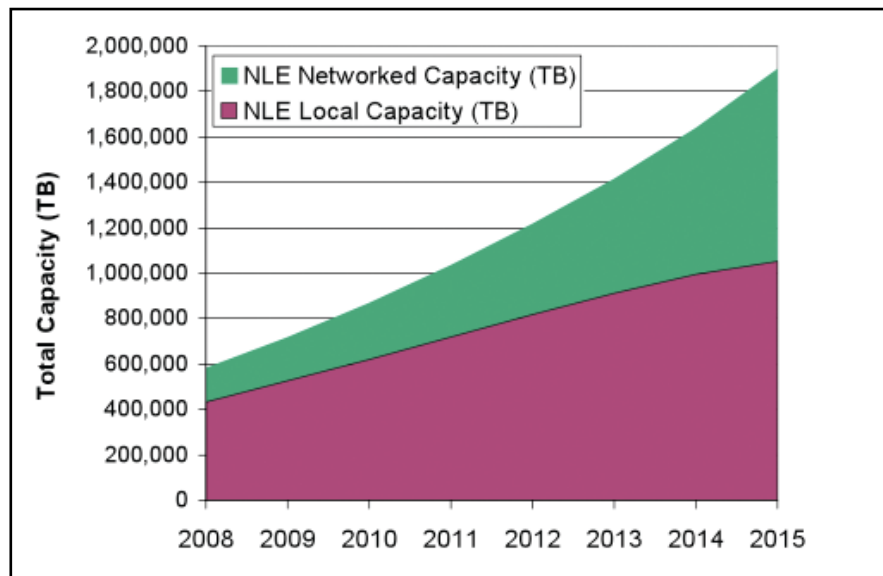
only the storage capacity that is needed for efficient post work is used. In larger facilities fast storage networks are making it possible to share high performance online storage resources between different workstations.

The vast bulk of working content in a post-production facility will be stored in a second and sometimes third tier of storage. The second tier for working content storage uses lower performing (and less expensive) storage systems to provide near-line access to content that is pulled into the online storage for delivery to the individual workstations as needed. These secondary storage systems are available on a storage network to the online storage tier using either block-based (SAN) or file-based (NAS) storage networking. If used to store all the material associated with a currently active post project we can call this second or near-line storage a working or active archive.

The third tier of storage in a post-production environment (if present) would use slower access storage devices and systems such as magnetic tape and optical disks. In this third tier of storage the storage media used may not be actively available on a storage network. This storage media must be loaded by people or by automation in order to access the stored content. The third tier can be used for backup and protection of the working archive. Backups and transfer of content on these media constitute so-called cold archives. These systems usually consist of the least expensive but the slowest source for retrieving content in a post-production house. Cold archives are typically used for preserving completed or non-active projects rather than currently active projects.

The chart, above, shows projections for the growth in annual demand for direct attached and network storage in professional post production facilities. We project much greater percentage growth of network storage in these environments,

Projections for Growth of Direct Attached and Local Storage vs. Network Storage in Professional Post Production Facilities



Source: 2010 Digital Storage for Media and Entertainment Report, Coughlin Associates

particularly in larger post houses, which drive the greatest storage volume in the market.

Network storage may use disk-drive arrays in a SAN environment, typically using high-speed Fibre Channel storage networks (particularly in larger post houses). The SANs may also include tape or some other lower cost, lower performance storage systems to provide data protection for the working archive.

For smaller shops, protection of the working archiving may be a simple process in which final cuts from working storage are written out in digital format to tape media and placed in a cold, offline archive along with the source tapes. Products used for data movement are typically asset managers but can also include backup applications run either automatically or as required.

For digital content written out from working storage, archive media uses commoditized IT storage like digital tape, optical storage, disk drives and even disk drive cartridges.

The final backup of a working archive in network storage can become the final cold archive for a project. Factors in the choice of cold archive storage formats include not only operating cost and cost per gigabyte, but the expected life of the archive media. Media life depends upon the quality of that media as well as the environment it is stored in. It is likely that

inexpensive off-the-shelf optical disk, for instance, may not provide long-term data preservation. Hard disk, while an easy to use format, may not be suitable for cold archives either, because of its possibly limited idle life span. For this reason, most cold archives use tape (e.g., LTO) due to its portability, quality, and long shelf life – upwards of 30 years. In addition, LTO tape drives have backward compatibility to two prior generations allowing read and rewrite capability to content saved on older tapes. In any event, media should be stored in a climate-controlled vault for high-quality preservation over a long period of time.

Larger organizations also use cold archives but in conjunction with some form of active archiving. In real-world deployments, active archiving tends to include a non-automatic hierarchical system consisting of multiple storage tiers.

A hard disk storage array may be used on the front end to provide faster content transfers between artists during ongoing and recently completed projects. Less frequently used files will reside on a local tape or even optical disk library system for near-line access.

In the case of post houses dealing with the artistic, budgetary and production needs of content creators, it is not uncommon for a production house to put post-production projects on hold. In this case the unfinished content needs to be put in a near-line storage environment where it can be kept economically until the production project can be started again.

Disk arrays used for active archiving typically use SATA hard drives since these provide higher storage capacity for a given price and provide good reliability. These arrays are used for less frequently accessed conditions that are expected in an archive application. There are even SATA storage arrays called MAIDs (massive array of inactive disks) that keep most of the SATA drives in the array powered down most of the time, providing significant power and cooling savings for a data center. MAID configurations are available from storage array vendors. Copan Systems, the initiator of MAID storage, was recently acquired by SGI.

Asset managers are used to manage the content on post-production storage systems and handle whatever level of automated tiering is possible. Some of the companies making software used for asset management and digital archive management in the entertainment and professional media markets (either by themselves or in com-

bination) include: Apple, Avid, Dalet, Etere, Front Porch Digital, Masstech Group, Quantum, SGI, SGL and XenData. Cold archiving is done by companies, such as Iron Mountain, that specialize in retaining records and assets in controlled environments.

Format Conversion and Long-term Archive Management

Over time, storage media formats and interfaces become obsolete. This poses special problems for a long-term digital archive. In the case of retained source media, extra copies of the reader devices for the media, as well as computer systems supporting these readers, are required to access the media. If the retention time is especially long, this could add up to a string of devices that can be used sequentially to convert older media content into more current formats with interfaces that more current devices support. This then becomes quite a collection of devices to maintain and a serious management issue.

Regardless of whether the original source media are retained or if the content data is transferred into a media archive system using hard disk drives, optical discs or magnetic tapes, obsolescence still forces the administrator to convert older storage formats into newer ones. Storage formats that include backwards compatibility mitigate this issue, but do not solve it. This issue becomes more pronounced as the size of the digital working volume or a digital archive increases.

In a sense, archiving (even cold archiving) should not be a static process. When the archive load becomes too large, choices will have to be made about which content to transfer and preserve on the new format. Format choices should always be made with a consideration of backwards compatibility. Otherwise archive transfers could become a constant process.

The rate of obsolescence for storage device formats varies with the devices. For hard disk drives the stability of the interface and the drive specification that controls that interface sets the backward compatibility. For SCSI drives and for Fibre Channel disk drives that use the SCSI specification there is very long format stability (backward compatibility) on drive commands.

On the other hand, the SCSI hard drive interface evolves to a faster version roughly every five

to seven years. ATA drives, previously used mostly in personal computers, now are commonly used in inexpensive storage arrays. SCSI and ATA made a conversion from parallel to serial connections a couple of years ago. The resulting SAS and SATA interfaces are completely different from the interfaces used in older format drives. Thus, disk drive interfaces change relatively frequently with time from an archivist point of view.

An archival system built around the use of disk drives must take this technology evolution into account and either have a means to migrate to newer drives using adapters as the old drives wear out or eventually move data from the older drive arrays to a new set. Therefore, the useful life of an active disk array is probably somewhere in the range of 5 to 7 years (the actual functional life with appropriate spares could be as high as 10 years). This time frame will also be influenced by service contract costs. After three to five years, service costs for disk drive arrays tend to go up dramatically, resulting in the desire to swap an array, even if it is functioning adequately.

For a MAID system with less active disks this time could be longer than 10 years. Often disk storage systems (as well as tape and optical) are replaced not because they no longer work but because there is technology with greater storage capacity, better performance and lower operating costs available.

The Popularity of LTO Tape in the Archive

The format obsolescence rate for tape varies depending on the format and technology. Different tape standards take different approaches. The most widely used professional digital tape format, LTO tape, follows the LTO Consortium's road map that communicates the members' intention to provide read/write capability one generation back and read capability for two generations back. The intent is to provide that capability at each specific generation's original density and performance. Thus an LTO-5 drive which is capable of recording 1,500 GB per cartridge at 140 MB/sec must be able to read and write on LTO-4 cartridges and must also be able to read LTO-3 cartridges, at the original specified performance.

The LTO Consortium has been introducing new generations every 20-30 months. Thus, archival

tape systems have a total format life (for reading) of about six years. However, an existing tape system could last much longer since the media is rated to last up to 30 years and tape storage systems (especially libraries) probably have a useful life of 10-15 years.

Optical storage media may be more stable in media format changes than tape but generally there is a major format change every 10 years and optical drive interfaces develop like those of other computer peripherals. Optical storage media of the right quality can probably last about 10 years under the right storage conditions and optical archive libraries probably have about a 10 year useful life. However, optical storage has its issues. Optical storage media has lower storage capacity per unit than tape and has a slower write time. In addition, labeling of optical media can be more difficult than tape unless the optical disks are contained in cartridges.

Thus an 8-to-10 year life of an archival storage system is likely for all these storage devices. Sometime before an older system is to be retired, the content it contains must be moved to replacement storage system.

To handle this format conversion, manual transfer processes could be used, but manual transfers are error-prone and time-consuming, taking staff away from other projects. Storage management software should include a way to automatically control the format refresh. Management software can read archive data from an older generation tape environment, verify the content is still intact, and then automatically rewrite to the latest generation of tape.

Migration of content between an older tape format to a newer one, such as, LTO-3 to LTO-5, can be done as a background operation and have minimal impact on the performance of production storage. On the other hand, migration from an older disk drive to a newer one sometimes cannot be done in the background and may interfere with ongoing production operations. In fact, migration of tape from one format to the next need not change the overall storage system footprint or automated process but only involves swapping out tape drives as required to convert and service a newer format. Since the storage capacity of the tape increases considerably between generations, the result is a significant change in overall storage capacity,

with little if any additional power required and no change to the storage system footprint. Therefore, tape migration may be a less painful path than replacing an older disk array.

By having this process follow rules defined by the user, the specific data retention and protection needs can easily be factored into the process. This also allows for maximum re-use of tape media and ultimately minimizes long-term data retention costs. Over a 100-year archive life, this process is likely to happen at least 10 times. This operation benefits from active software management to make sure such media progressions occur according to a pre-set schedule and that transfers are successful before the older media and storage systems are retired.

Creating Better Metadata for Post Production

Creating better metadata to represent the data in production storage systems and creating the means to manage this metadata are keys to better search and discovery of historical content. While today much metadata is entered manually, technologies are being developed to create automated metadata generation from the original content and to enable searching data based upon matches with audio, still or moving video content. As these technologies develop, the resulting metadata could be incorporated into professional video metadata formats such as MXF to create powerful ways to access older content archives.

Storage devices and storage management software should be designed to make use of these automated metadata generation capabilities as they are developed. Production storage management software should also be able to search through active storage as well as archived content, creating new metadata based upon rich media attributes. In the absence of this, or in environments that use customized metadata formats, the storage management software should be as tightly aligned with the asset manager as possible, so that artists can track and manage content as they need.

Decreasing Overall Costs of Production Storage

The cost of production storage is a function of several factors. Some of the more important of these factors are storage media utilization and

the costs of maintaining the data center where the storage assets are kept. Storage management software is a key element of reducing the overall cost of storage.

Although generally higher in archive environments, storage capacity utilization is generally rather low. Disk storage utilization of less than 50 percent is not uncommon. It is possible to improve storage utilization by creating virtualization of

By reducing the overall hardware and facilities costs (which tend to be much greater over the long run), good storage management software can more than pay for itself.

storage assets so that storage can be provided to the user as needed rather than relying on pre-provisioning storage. By reducing the overall hardware and facilities costs (which tend to be much greater over the long run), good storage management software can more than pay for itself.

Storage management that can find duplicate data and not backup multiple copies of the same file is one method to reduce overall storage capacity, as well as reducing the bandwidth requirements for moving this data around.

The cost of maintaining storage facilities greatly exceeds the initial hardware costs. Data center power, cooling, floor space and environmental issues are essential to successful operation of an active storage facility. Reducing the heat load from the storage systems will reduce overall facility costs. The use of smaller disk drives in drive arrays (particularly if these are lower-RPM SATA drives) will reduce the heat generated by these drives, and since the smaller drives can be packed tighter in an array than larger disk drives, the overall storage footprint can be smaller for an equivalent storage capacity. This reduces the size of the data center space required for the active archive.

For larger archives, tape is an important format because it does not require a large amount of power or floor footprint. Media is offline except when in use in a tape drive. The library itself requires no power except for the robotics and drives, which are often idle. Tape libraries can scale to multiple petabytes of storage without requiring significant power consumption in comparison to a disk array of the same capacity.

Broader Asset Management Systems

A final topic of discussion is how production storage and archiving fits into a larger asset management system. An asset management system is a set of processes designed to control digital assets used in active and inactive projects. These processes should cover every aspect of the content lifecycle including tracking metadata, the data's physical location from active usage to archive, and the overall handling of data protection.

While there is no single, catch-all product that covers all aspects of asset management, some broad strokes can be made about the design of a system. An asset management system should do the following:

- *Protect the digital assets for many years through active management and checking of the hardware and storage environment.*
- *Perform systematic format updates and transfers of the digital content as needed to avoid potential access loss due to format obsolescence.*
- *Alert the administrator of problems with the content, environment or storage hardware early enough for the administrator to take action and if necessary shut down and protect assets in the event of a crisis.*
- *Provide organized access to the digital archives by retaining an appropriate database of content metadata and perhaps indexing of the actual content on those assets.*
- *For an active archive, control the access and delivery of digital content when needed and in a timely manner.*

- *For a cold archive, initiate delivery and mounting of the media from the cold archive to where it is needed.*
- *Manage disaster recovery requirements such as mirroring or backup of managed data to remote data centers to make sure content can be recovered if the original copy is somehow damaged.*

New developments in long-term archive management include:

- *Active indexing of still and moving digital images during the archiving process—more advanced metadata creation and management—allowing easier search and use of this content.*
- *More active and continuous checking of archived content to verify data integrity and to detect growing problems.*
- *File-based access to digital content from the storage media to enhance content access.*

Conclusions

Management software and hardware play an important role in post-production. The decreasing cost of digital storage and the capture and editing of content enabled by digital technology has increased the number of facilities and producers of content. Although the recession of 2008 and early 2009 impacted post-production spending as it did other storage markets, the current market recovery and moves to higher resolution stereoscopic content are driving new demand for post-production digital storage. Improvements in the economy and efficiency of storage systems is also helping to reduce the capital and operating costs of M&E storage.

Digital technology and increased available communication bandwidth have also allowed the development of new ways to distribute content, such as the Internet and mobile phone networks. This has also increased the supply of new professional (as well as non-professional) content. Higher resolution as well as stereoscopic content is being required in professional media and the amount of digital footage being acquired for the final produced product has increased considerably. In addition, older analog content is being digitized in greater amounts with time. All of these factors

have caused digital storage capacity requirements to swell.

As total content capacity increases, so does the amount of content that is being archived in either cold or active archives (or both). Long-term preservation of large digital archives will lead the industry to solve new issues associated with format conversion, metadata creation and management as well as methods to reduce the total cost of operating a digital media archive. Preservation of digital data from the possible destruction of the primary storage system can be dealt with by having off-site copies of the content or by maintaining remote mirrors or backup of the content. Migration of storage formats may be easier for tape than for drive systems for long-term data archives. ■

Sources:

- Digital Storage in the Post Production Market, 2008, Coughlin Associates
- 2010 Digital Storage For Media and Entertainment Report, Coughlin Associates

Thomas M. Coughlin, president of Coughlin Associates, has been working for over 30 years in the data storage industry as an engineer and senior level manager. He has over 60 publications and six patents to his credit. He is active with IDEMA, the IEEE, SMPTE, SNIA and other professional organizations and is the founder and organizer of the Annual Storage Visions Conference and a partner to the annual Consumer Electronics Show as well as the annual Creative Storage Conference. Coughlin Associates provides market and technology analysis on digital storage and applications as well as Digital Storage Technical Consulting services. His publications include an annual report on data storage for professional media and entertainment and Digital Storage in Consumer Electronics: The Essential Guide, published by Newnes Press in 2008. He also publishes the quarterly The Data Storage Technology Newsletter. For more information go to www.tomcoughlin.com. Contact him at 408-978-8184 or tom@tomcoughlin.com.

Just about every new camera system has ditched the tired, old idea of recording to tape. But not all tape needs to go the way of the Dodo. LTO-4 is still the best option for backing up and archiving the footage you've shot on your tapeless camera.

2

Archiving in Action

The DIY Archive for Independents

Create a Secure, Low-Budget Archive Using LTO-4 Tape on the Mac

by Helmut Kobler

Just about every new camera system has ditched the tired, old idea of recording to tape. I thank my personal god for that, but not all tape needs to go the way of the Dodo. It turns out that a particular tape format called LTO-4 is still the best option for backing up and archiving the footage you've shot on your tapeless camera (and no, you shouldn't archive valuable footage to spare hard drives that sit on a shelf—drives left static for a year or two become unreliable).

LTO tape (stands for Linear Tape Open) has been around for many years, and is already used by financial companies and government to keep track of your old banking statements and parking tickets. LTO-4 is the newest flavor of LTO, and its benefits are considerable:

- ▶ A single \$130 - \$150 LTO-4 tape (street prices often much lower) can store 800 GB of uncompressed video—that's 26 whopping hours of 720/24p video from a Panasonic P2 camera, 38 hours of 1080/24p footage from a Sony EX3, or nearly 8 hours of 4K footage from a RED camera.
- ▶ An LTO-4 tape lasts about 30 years in normal storage conditions (i.e., don't leave it in your trunk indefinitely).
- ▶ Some LTO-4 drives can read/write data up to 120 MBps—i.e., much faster than reading/writing to a conventional SATA hard drive—so incremental backups go in a hurry.

There is, however, a problem with LTO-4 tape... and it's not really the cost. In fact, a fast, reliable LTO-4 system can be had for under \$3,300, which is reasonable for most professionals. LTO's real problem has been the confusing and risky process of finding the right LTO-4 drive, the right interface card, and the right backup software, and see them all work reliably together.

That's what this article is for. After months of research, I've found a Mac-based LTO-4 solution that's working great on a single workstation. You can buy it yourself piecemeal, or buy it as a bundle from one company (TOLIS Group, www.tolisgroup.com).

Here's how the various components work together.

The Tape Drive

Many companies make LTO-4 drives, but I chose the HP Ultrium 1840, which has a list price of \$3,999 (naturally, it can be found at online retailers for much cheaper). The 1840 is a "full-height" external drive that's about 12.5" inches long, 9" wide and 5" tall, and can easily fit on your desktop. It's a speed-demon as far as LTO-4 drives go, delivering up to about 120MB/s of real-world throughput. And it's got a confidence-inspiring 3-year warranty.

HP sells the 1840 with either an Ultra320 SCSI or SAS interface, but I chose the SAS option because it delivers top speed through a single cable, without any configuration headaches that I associate with other SCSI formats.

The 1840 is as simple to use as a Fisher Price toy—the front of the unit sports a Power button, a Tape Eject button, and a couple of status LEDs. All you really have to do is gently nudge your LTO tape into the Ultrium, and the drive takes it in.

When the Ultrium is on, you can definitely hear its fan. I'd say it's about as noisy as a typical 8-drive RAID, and definitely more noisy than a Mac Pro tower. When the drive is searching the tape, there's also a little robotic whir thrown in for good measure. Given the noise, the 1840 isn't something I want running all the time on my desk. That's okay, though, because I can just turn it on when needed, and my backup software finds it immediately—no need to reboot the computer.



Note: HP also sells a smaller, less expensive “half height” drive called the Ultrium 1760 (\$2,749), but I chose the 1840 (above) because its max speed is about 20MB/s faster, and it runs about 12 db quieter. If cost is a major concern, though, the 1760 is still a great option.

The Interface Card

The Ultrium drive connects to a computer via a Mini-SAS cable, so I installed an ATTO Technology H380 card (\$395) into my Mac Pro to gain two SAS ports.

The H380 is a no-hassle kind of card. You do have to install a driver to use it (unless you’re running the new Snow Leopard OS, which comes with drivers built in), but the card runs cool, has no fan, and doesn’t interfere with the Mac’s ability to sleep. Although the card is designed for an 8x slot, you can install it in a 4x slot if it’s just running a single LTO-4 drive.

If your Mac is running low on free card slots, ATTO also sells a highly-rated RAID card called the R380 (\$1,095), which includes two SAS ports that can drive either a 4- or 8-drive RAID and an LTO-4 drive simultaneously, at high speeds. The R380 is a great way to make the most of the Mac Pro’s rather paltry three spare expansion slots.

One more thing: most LTO-4 drives use SAS or SCSI interfaces, but a few drives, such as CacheA’s PrimeCache (\$7,995), can connect to your computer or network via an everyday Gigabit Ethernet cable. The drawback of an Ethernet connection is that it’s only half as fast as SAS, but that may be worth it if you want to archive video from multiple workstations. Since I’m archiving video from a single Mac, I opted for a zippy SAS-connected LTO drive.

The Backup Software

The final ingredient in my LTO-4 cocktail is backup software called BRU Producer’s Edition (\$499), from the TOLIS Group. The Mac operating system (Windows, too, for that matter) doesn’t work with tape drives natively, so I need BRU to control the HP Ultrium and read/write data to its tape.

BRU is a full-featured backup application, but its beauty is its simple QuickArchive mode, which provides a clean, streamlined interface for writing data to tape. To get started, just drag any files/folders from the Mac’s Finder into BRU’s Archive window, and then type in a name for your Archive session—for instance, “Wrestling B-Roll.” Then just hit BRU’s Archive button, and watch it go work writing the data to tape.

Writing goes relatively quickly—again, the Ultrium 1840/ATTO H380 tag-team manages a real-world data rate of about 120MB/s, but BRU also needs time to verify that your data has been accurately written to the tape. The bad news is: the verification pass usually takes as much time as the copy itself. The good news is: the wait is still tolerable. For instance, backing up 243 GB of P2 video took an hour and four minutes—not too shabby for ultimate, peace-of-mind security!

To restore an archive session, just click BRU’s Restore button, which brings up a new interface that lists all the archived sessions you’ve created. You can restore a full session or pick and choose files/folders from that session. You can also type in a search string and let BRU show you all the backup tapes you’ve made that include similar items. When you’ve selected your data, another button click lets you choose where it will go—most likely, you’ll save it your editing RAID or a spare hard drive. Finally, click BRU’s Restore button, and BRU starts reading your tape (reading goes just slightly faster than writing, including verification).

Go beyond BRU’s QuickArchive mode and you’ll find plenty of advanced features, such as the ability to schedule automatic back-ups for certain times, or to back-up only certain files from certain dates, etc. That’s all welcomed, but will probably be overkill for many shooters/editors who will simply connect their latest batch of footage-carrying cards or drives to their Mac, and back up immediately.

There is one potential gotcha when working with BRU, and that’s that no other backup soft-

ware can recognize your BRU-formatted tapes. 10, 15, 20 years down the road, if you want to restore some old tape in your archive, you'll need BRU-compatible software to do it, not to mention a tape drive that reads your LTO-4 tapes. The just-released LTO-5 and upcoming LTO-6 drives, which are on the drawing board in labs today, are guaranteed to work with LTO-4 tape. But who knows after that.

Given LTO's long-term orientation, you might wonder if TOLIS Group will be around to sell a version of BRU that works with whatever newfangled hardware and OS you'll have in the future. There are no guarantees, of course, but TOLIS and BRU have been around since 1985, and it's TOLIS policy to keep data compatible year after year. I definitely felt reassured to learn that today's BRU can still read data files written more than 20 years ago, when BRU ran on Amiga computers and backed up files to floppy disks and prehistoric tape.

TOLIS has also placed BRU's source code with a number of long-time clients and media industry groups—for instance, the code is currently with NARAS (National Academy of Recording Arts and Sciences, which runs the Grammy Awards), and TOLIS is working to place BRU code with SMPTE and the Society of Composers & Lyricists as well. In other words, if something ever happens to TOLIS Group, BRU can still rise from the ashes.

Bringing It All Together

You can build an LTO solution with other tape drives, interface cards and backup software, but a

quick check of online support forums shows that some combinations don't work well together. On the other hand, the HP Ultrium 1840, ATTO's H380 card, and TOLIS' BRU Producer's Edition work like a well-oiled machine, so if you want a hassle-free experience, I heartily recommend this trio.

You can assemble the LTO hardware by scouring the net for cheap-o deals (including grey-market hardware); if so, you might be able to acquire everything you need for as low as \$3,300. A more trouble-free approach would be to consider buying everything in a bundle from TOLIS Group. The 1840 drive, interface card, SAS cable, BRU software, a blank LTO-4 tape, a drive cleaning tape, and 12 months of telephone support go for \$5,180 (a cheaper bundle based on HP's smaller 1760 LTO drive is \$4,015). Opting for the bundle means paying just about list price for everything, but you'll get it from one source, and all the support you need to get your feet wet. ■

StudioDaily contributor Helmut Kobler, a former video game designer, is the owner of K2 Films, a small production company in Los Angeles that makes one independent film at a time. He is the author of several Final Cut Pro for Dummies books, as well as Anatomy of a Guerrilla Film: The Making of RADIUS, which details his experience as a first-time director.

Digitizing a Museum Tape Archive

How the Nonprofit HistoryMakers Transformed 2,000 Oral Histories into a Digital Archive for Generations to Come

By Steve Kwartek

What does it take to digitize the nation's largest African-American collection of video oral histories — some 2,000 irreplaceable interviews with significant African Americans from all walks of life?

Two graduate student assistants, a Front Porch Digital SAMMA Solo system, and a whole lot of advance planning.

The collection belongs to The HistoryMakers, a Chicago-based nonprofit organized in 1999 to raise awareness of the contributions African Americans have made to American society and culture. The organization's purpose is to produce and make available to the widest possible audience video interviews with African Americans who have made a difference, whether they are well-known or unsung. The oldest person represented in the archive is 106 years old, and the youngest is 29.

With the election of Barack Obama as the first African-American president, The HistoryMakers' priceless testimonies are more relevant today than ever before, according to Julieanna Richardson, the organization's founder and executive director. "We have the only video oral history interview available of President Barack Obama as an Illinois State Senator. We receive almost daily requests for access to our collection," she says.

The HistoryMakers' archive includes footage shot in a variety of formats, including ¾-inch U-matic, VHS, Digibeta, and DVCAM. But the overwhelming majority is on Beta SP videotape. Betacam SP was developed in 1986 and remained the industry standard for most television production through the late 1990s. While the image quality is excellent, the medium deteriorates over time, even when it is stored in optimal conditions. Since summer 2009, The HistoryMakers has been in the process of migrating

its content, some 7,000 hours of it, to digital storage where it can be preserved for the long term and also made more readily accessible to more people.

For all its ambitions and accomplishments, The HistoryMakers is a small organization, with a staff of only four people, including Richardson, and an annual budget of about \$1.2 million. In 2007, The HistoryMakers forged a three-year collaboration with the University of Illinois. This collaboration provided The HistoryMakers with fully paid graduate assistants and a University archivist who lent her expertise and experience with established archival principles.

Understanding Digital Storage

According to Richardson, digital storage is initially less intuitive than analog storage. Anyone can understand a traditional videotape with a label on it on a shelf. An LTO data tape with a barcode on it that can be read only by a machine, and with information including many layers of metadata, initially seems opaque. To make sure all that information is comprehensive



Daniel Johnson, Archive Assistant

MUSEUM ARCHIVING AT A GLANCE

TARGET FACILITY: Museum, corporate or nonprofit organization with substantial video archive that needs digitizing and archiving

WHO: The HistoryMakers, a Chicago-based museum containing the country's largest video collection of African-American oral history

STAFF: 4 full-time employees

ANNUAL BUDGET: \$1.2 million

CHALLENGE: Digitize and archive a 7,000-hour video library of interviews collected on a range of analog and digital formats

ARCHIVING SOLUTION: Front Porch Digital SAMMA Solo System, which digitizes all varieties of analog tapes to LTO4 Tape in multiple formats (including JPEG2000)

COST: Contact Front Porch Digital for more information (www.fpdigital.com).

ADVANTAGES: Fast, efficient and easy-to-use, the SAMMA Solo also features a unique sensor and cleaning process that flags failed tapes to be put aside for evaluation and removes any deteriorating oxide, metal particles, or dirt from incoming tapes for high-quality transfers.

DISADVANTAGES: Depending on the size of your archive—and staff—you'll need at least six months to a year of advance planning to organize your corresponding database catalog before you begin digitizing, which can take even longer.

WHAT YOU NEED: Any analog tape library that needs digitizing, whether VHS, Umatic, Beta SP, Digibeta or a combination of the above; Windows XP; a long lead time (plus an eager intern, graduate student or entry-level staffer)

and useful, the planning and cataloging process that leads up to digitization has to be done right.

At The HistoryMakers, preparing for the migration required a full year of methodical organization and assessment in order to ensure adherence to the profession's best practices. Among the steps accomplished during the year were digitization of The HistoryMakers' paper records and the assignment of accession numbers to the interviews — in other words, much of the back-end work of the database. In addition, the staff identified which tapes they wanted to digitize first and evaluated the formats on which the data was stored as well as the type of digital formats they wanted to use for storage for preservation purposes.

"Because the assessment of the collection and its reorganization took place before the SAMMA Solo system was implemented, the project has gone extremely smoothly," says Richardson. "To be candid, it also helps that our collection is only

nine years old, so it's possible for us to have quite a comprehensive knowledge of it."

Richardson considered all possible means of actually performing the migration to digital but soon recognized that only one was practical. Front Porch Digital's SAMMA Solo is the only commercially available system that provides fast (real-time), semi-automatic migration of content from analog videotape to digital storage on managed LTO data tape.

"I like two things about it especially," Richardson said. "First, it is very efficient and easy to use, which is critical given our staff and resource pool. Second, it was developed by people who understand not only technology but also archiving processes. It is truly a unique product."

The SAMMA Solo simultaneously ingests and encodes content stored on The HistoryMakers' Betacam SP tapes into multiple file formats, including JPEG 2000, with the digitized recordings incorporating metadata describing the content's



Rows of Archival Boxes and Beta Tapes.

condition. As a result, a history of the tapes' condition is preserved along with the digital file.

An important component of the SAMMA workflow is the videotape cleaner, which helps assure the best possible quality when the tape is played back. The cleaning process removes any deteriorating oxide, metal particles, or dirt from the tape and prevents head clogs, which can keep the video and audio recording from playing out correctly. The cleaners also have a sensor alert system to flag failed tapes so that they can be put aside for evaluation and repair. The cleaner allows a high-quality transfer to be made even from tapes in less than perfect condition.

Creating a Timeline

The HistoryMakers' digital migration project began in the summer of 2009. Beginning last fall and continuing through the spring semester of 2010, two graduate assistants from the University of Illinois' school of information and library science work consecutive shifts from 9 a.m. until midnight five days a week. The migration process begins with the assistant inspecting each tape and placing an identifying barcode label on its cassette spine. A videotape that passes inspection is loaded into the SAMMA Solo, and the migration begins with the robot scanning its label for identification. Each tape is then moved by the robotic arm to the videotape cleaner. Once it is clean, it is moved to the VTR for playback, with high-resolution content converted to files automatically moved to high-density LTO-4 data tapes—all within the automated system.

"The SAMMA device automates repetitive processes and evaluates and monitors tape quality using signal analysis, so most potential problems are identified and

corrected without our having to spend staff time on them," says Richardson.

The HistoryMakers launched the project with about 30 interviews taped in 2008 deemed valuable but which had quality problems preventing them from being easily dubbed. These included one recording with the former governor of Virginia, Douglas Wilder. If the project is able to continue at its present rate, the staff anticipates the entire collection can be migrated in two-and-a-half years.

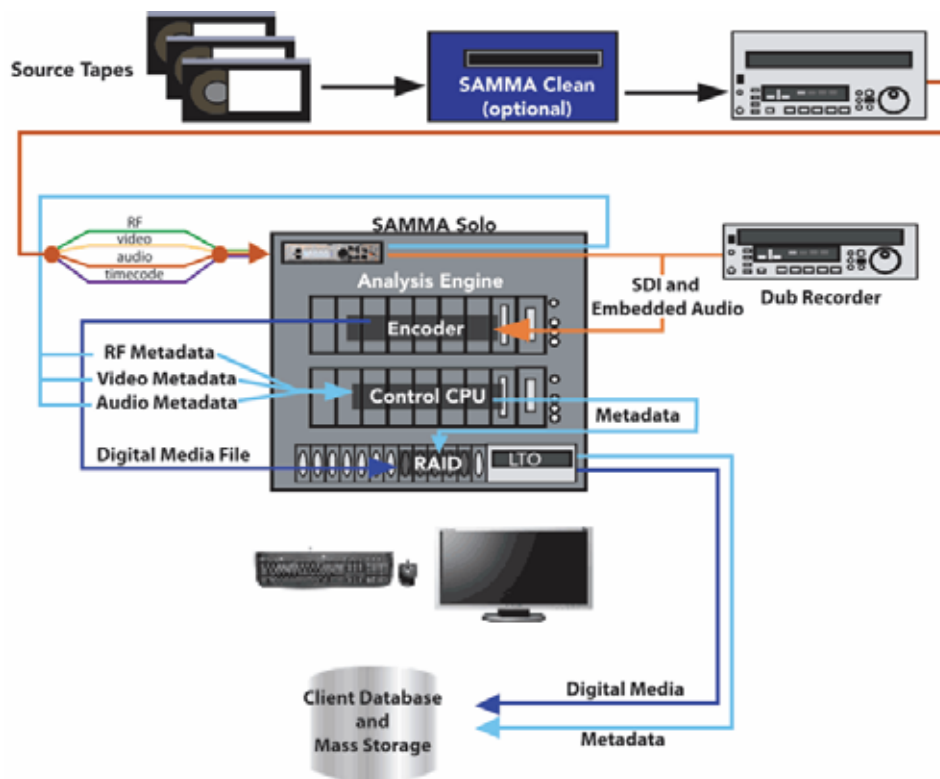
While the organization encountered fewer problems than expected once the migration began, there have been some lessons along the way. For example, the

staff recognized early on that they needed to institute a final step to quality check the JPEG 2000 files and their associated metadata. Likewise, Richardson says she now sees that it will be valuable for future keepers and users of the collection to have documentation of the decision-making process behind its cataloging schema. Finally, The HistoryMakers' current MAM (Media Asset Management) system is based on Filemaker, and it needs to be upgraded.

Richardson describes the decision-making process ahead of implementation of the SAMMA Solo as "agonizing" because she had to educate herself about the technical side of archiving, including digital formats. However, that investment of time and energy will pay off in the long run. For example, The HistoryMakers previously has had to turn down offers of video collections if they weren't Beta SP-based. Now, the flexibility of the SAMMA system will enable growth by acquisition of collections stored in a much wider variety of formats.

"Being able to grow our collection by acquisition, something we could not have done in the past, is huge," she says.

To make the collection even more useful to future scholars and the general public, The HistoryMakers is also working with Carnegie Mellon University's Informedia Digital Video Library group to take advantage of improved search algorithms that rely on artificial intelligence to more nearly approximate the way human beings actually look for information. For example, these algorithms will make it possible for someone searching the video archive for information on the Great Migration to find a video in which the subject mentions that his



or her family moved from Mississippi to Chicago in the 1930s — even though the testimony does not include the phrase, “Great Migration.”

Like other oral history collections, The HistoryMakers collection recalls the earliest efforts to capture the voice of a people, but at the same time it is distinctive in that it introduces the modern advantages of technology for preservation and increased public and scholarly access.

“Most people don’t really recognize what an archive is,” says Richardson. “You say the word, and they think of something old and dusty. But our video archive is not; it’s alive and will be very easy to use and access. I have sought out the best of the best to help, and as a result, we are at the leading edge of what is going on in the field.”

Richardson is a Harvard-trained lawyer who majored in American studies and theatre arts as an undergraduate. She believes that storytelling is a natural part of who she is. When she founded The HistoryMakers, it was with the goal of leaving a legacy that would incorporate African-American history and culture into the mainstream of American society. A key step toward that objective is to make the material available for mass distribution—on television, for example. Digitized

information can be accessed via the Web, and it can be duplicated easily and automatically, making this current migration project an integral one in furthering the mission.

“The SAMMA acquisition represents a very significant step for us in the history of our project,” says Richardson. “It secures our future and the preservation of this material. You don’t work as hard as we have to create something, only to discover it’s lost because it wasn’t done correctly.” ■

Steve Kwartek is North American regional sales manager, nonprofit organizations, for Front Porch Digital. He is responsible for accounts including the U.S. Library of Congress, the National Archives, The Shoah Foundation, and the United Nations. An authority on media access and preservation, he has worked with the SAMMA semi-automated digital migration system since its early development stages, in both sales and marketing capacities, and since 2009, as part of Front Porch Digital. Kwartek came to SAMMA from VidiPax, a leading videotape restoration company where he served as a restoration specialist and sales and marketing manager. He began his career as a production manager at the Museum of Television and Radio (now called the Paley Center for Media) in New York City.

Archiving on Set: Production

Getting The Group Together: Managing 8 Sony EX3 Cameras, 8 Mac Laptops, 20 Flip Ultra HDs and 500 Hours of Footage for a Reality Show Pilot

By Beth Marchant

In the summer of 2009, producer David Sutcliffe took a television production crew and 20 other willing participants to a ranch in Laurel Springs, California, to tape a grueling, six-day group therapy session based on the Core Energetics method. The result is a reality show-cum-documentary pilot with which Sutcliffe, an actor known for his work on *CSI* and *The Gilmore Girls*, has an unusually personal connection. He says he was transformed by the therapy and as a producer, felt that the method lent itself to the reality show formula of trial and transformation. “The (Core Energetics) work is very dramatic,” he says, producing “very cathartic emotions.” He is currently working with an editor to figure out the episode format while he looks for a home for the series.

Brought into the project by DP Jesse Feldman, Hanton’s long-time friend, DIT Emile Hanton supervised production during the shoot, and with the help of reseller and integrator Media Distributors, came up with an on-set workflow that allowed the crew to loop in multiple cameras and hours upon hours of footage—more than 500 hours in all. Eight main Sony PMW-EX3 cameras captured the primary action and 20 Flip Ultra HD cameras, given to each participant to record their innermost thoughts in any way they saw fit, recorded everything else.

So how exactly do you manage all that footage without disrupting the natural flow of the sessions themselves? Hanton says Media Distributors’ Archive Station instantly archived and backed up to a SAN and LTO drives. The company’s Constellation VCM, built into the turn

key, single train case-contained system designed by Media Distributors’ Nathan Adams, let Hanton and the crew search, retrieve, back-up and archive

all the footage that was downloaded from the cameras onto eight Mac laptops.

“The Archive Station, essentially, with its 12 TB RAID, networked the laptops via Gigabit Ethernet and fiber to a Facilis 24D SAN,” says Hanton. “The MacPro in the system handled on-set editing and file conversion. After we backed up the original card structure, we transcoded everything to editable QuickTime, so it was ready to go when editorial needed it. Overnight, the day’s material would get backed up to LTO tape automatically.” Media Distributors also rents and sells a Blu-ray Disc version of the Archive Station.

It’s About the Journey

“There is no linear story, per se, to this show,” says Hanton. “There is no cash prize at the end. It is about the emotional struggles and personal discoveries of these people. David, the director, and Jesse, the DP, were looking at possibly trying to get the Sony EX3 cameras into the participants’ rooms, and then someone brought up the idea of the Flip cameras as a great alternative to just a traditional, sit-down confession-cam. It also solved the problem of trying to get cameras in the rooms of these participants, without having to worry about all the wiring that goes along with it. They wanted to move beyond the sweeping camera moves and the calculated, big reality show style that has been seen so much.”

The addition of the Flip cameras proved invaluable, although the format added another layer of technical hurdles. “We knew from the start that those clips were going to look different than the EX3 footage,” says Hanton. “But that was more of a benefit than anything else, because we wanted to do more than the usual confession-cam shots, even though on other reality shows, you often get a DV format to sig-

Tips for Better Archiving, from Set to Post

Media Distributors' Nathan Adams on what to do when backing up footage in real time



According to on-set archiving expert and former post supervisor Nathan Adams, “You can never have too many back ups.” His on-set advising and system integration work for Media Distributors, a rental house that also designs and integrates storage, archival and asset management workflows and systems, means Adams has heard countless horror stories about disappearing and corrupt footage, especially as more facilities transition to tapeless workflows and use solid-state media on a regular basis.

Beyond investing in or renting asset management systems and specific archive appliances (Media Distributors makes the custom Archive Station, which includes Constellation VCM), Adams points to several rules of thumb during log and capture that will help you clear a path to both easier retrieval and secure, long-term storage. “First and foremost, be organized. Always know where your project is going and work backwards to plan your path,” he says. You can’t do that, he says, without a consistent naming convention across all your media. “Not only will being organized help you keep track of your media, it will keep you employed.” And don’t forget to bill your client for the daily, project and long-term backups you routinely make.

What advice does he have for those recording to reusable flash drives? “You want to always reuse your media because it will save you a boatload of money.” You just have to be smarter about how you offload them, he says, to avoid accidental erasure or overwritten files. “And don’t forget: a FireWire drive is definitely not for long-term archiving,” he cautions. “It is a common default but it’s not a good archiving choice. If the drives break down, your originals are lost for good.”

When working with Sony XDCAM EX SxS cards and Panasonic P2 cards, Adams offers this advice:

1. Always keep the footage saved in multiple locations
2. Check the footage after a copy to ensure there are no corruptions
3. Copy off the SxS or P2 card over USB 2.0 at 2GB/min



When it comes to RED CF Cards, however, he’s a bit more cautious. “Read the disclaimer first: remember, you’re a beta tester,” he says.

Determining how much on-set or in-studio storage you need comes down to knowing your footage’s frame size, frame rate, color-bit depth, color sampling, aspect ratio and data rate. “Because each codec has a specific data rate (MB/s) for each specific frame size and frame rate, once you know what you’re working with, you can calculate how much storage you need and how fast you need it to move your files there,” he says. If you need help figuring that out, he recommends AJA’s Data Rate Calculator (<http://www.aja.com/products/software/>).

For long-term back ups, Adams recommends both LTO tape and Blu-ray. “LTO-4 costs only about \$150 for 800GB, which is much cheaper than Blu-ray (about \$1 per GB). But Blu-ray is still 10 or more years of secure storage for your masters,” he says. —*B.M.*



nify the shift to the one-on-one confessional. The fact that it was a lower-quality compression was fine. We were surprised, however, at what we ended up with.”

What they got was a broad range of footage with very personal, diary-like video entries. “The material that we got was absolutely incredible,” he says. “And we know that’s because our subjects were freed from any kind of schedule or requirement, like ‘Go to this particular room and tape your confession.’ We were worried that if we had only a few confession cameras, with 20 participants, some might overlap—we’d introduce another stressor. With the pocket-size Flip cameras, they could take them anywhere they wanted to on the entire campus, a beautiful ranch near Santa Barbara. Instead of the same shot of everyone sitting in a chair with a camera, we got people to open up for 30 minutes to an hour, sitting in a secluded but often beautiful spot, opening up about everything that they were going through during the process and in their lives. The footage just adds such a personal dimension to the project.”

The Flip cameras had four dedicated laptops for downloading and transcoding footage, says Hanton. “We decided to convert all the Flip video, through Apple Compressor, to ProRes. We knew that ProRes would be our final form, where we’d do color correction. So it made sense to do this right away. When they came back from the participants, we’d just plug in the USB and download over GigE to the RAID and the NAS. We processed about 200 GB, or 50-to-75 hours of footage. They also served as gifts to the participants, so they took them home with them. Some of them are still recording, after the fact, and David wants to use as much of that as he can.”

Wireless transmission was another way the crew kept its footprint on set to a minimum. “On the first day, we had all of our cameras rolling on the participants arriving,” Hanton says. “We spent four days on the group therapy sessions themselves, with a morning session and an evening session, each three hours long. We only had three camera operators in the room at one time to film the sessions. We’d rotate out the operators every 45 min-

utes. Since we knew we’d be juggling cameras, and we wanted everyone, including the subjects, to be able to move freely, we settled on the IDX CAM-WAVE to capture the full HD signal, via HD-SDI out, from the EX3s wirelessly in our video village. This also let David and Jesse view it in full HD.”

If he has to name the most challenging part of the project, Hanton says it was pre-production legwork. “How quick can we turn around cards, and how fast can we download footage, how many laptops do we need—doing the leg work up front to figure these questions out was the hardest part. Once we got the system in place, however, it was flawless. It was so fast, in fact, that I realized in the end that we could probably have done with 15 less cards because we turned things around so quickly. On set, I manage everything myself, so I couldn’t afford for the workflow to break down. Nathan and everybody at Media Distributors took the time to dial up the system before we got to the shoot, and it made all the difference.” ■

Beth Marchant is the founding editor-in-chief of Studio/monthly and the co-editor of StudioDaily.com, a collection of Web sites, including Studiomonthly.com and filmandvideomagazine.com, dedicated to production and post-production technology. She has spoken on a range of digital content creation topics at the National Association of Broadcasters convention, the Film Finance Forum and at the Writers Guild of America. Her articles and photographs have appeared in Yankee, Harvard Magazine, on the syndicated Gannett News Service and in numerous Studio Group publications in print and online. She also reviews consumer technology for the site Techlicious.

Archiving on Set: Post

Authentic Entertainment's Archiving Workflow for Naughty Kitchen and Ace of Cakes

By Jessica Clegg

Authentic Entertainment, a production and post facility located in Burbank, CA, works primarily with reality shows including *Naughty Kitchen* and *Ace of Cakes*. This highly edible entertainment requires many hours of shooting footage and a hardy, reliable back-up system. Will Piszieski, post producer, explains how Authentic has streamlined its fast-paced workflow with a Cache-A Prime-Cache archive appliance.

How long have you been using the Cache-A Prime-Cache archive appliance?

I believe we were the first production company to get it back in June 2009 when it first came out. I was very impressed by it at NAB. We were doing a new show for Oxygen called *The Naughty Kitchen*, which we wanted to shoot completely file-based and we needed some kind of long storage solution we could sell the network on. At the time, sending FireWire drives to the network was not an acceptable option.

So how do you send files with the Cache-A?

The Prime-Cache archives to LTO tapes using a format called tar, which has been around for a long time. Basically you can put any Cache-A written tape in any standard LTO-4 drive and read it using tar.

Could you describe the installation process? How was it adjusting to the new system?

The installation process was very simple. It was actually just myself and our IT supervisor. We utilized a CAT-6 Ethernet cable and connected it to our local network. After doing this, we could access the system from any of our 35 edit bays and eight After Effects stations. That means any

of my group can archive footage whenever they need to do so.

How's the interface?

It has a very user-friendly interface. It keeps a database that tracks every tape that you've made. Let's say hypothetically you've got media on a FireWire drive and that FireWire fails, it's very easy to go into the Prime-Cache's interface and locate a tape that you made as a back-up and pull it in the recovery process.

Talk a little bit about your workflow using *Naughty Kitchen* as an example.

We have a digital technician who works in Dallas where the show is shot. He's given these SxS media cards, which are shot using Sony's EX3. He offloads those cards onto a primary and a back-up drive, and then ships the FireWire drive back to Authentic every couple days. That FireWire contains just the original MXF files that are recorded on the card. Our assistant editor then begins the process of unwrapping those files, turning them into XDCAM QuickTime files we would use in Final Cut Pro, then takes the MXF files and drops them onto the internal hard drive of the Cache-A device, and then begins to archive them onto LTO tapes. We choose to archive each day on its own tape, even though we probably don't fill up any of the tapes—the tapes hold about 800 GB—but we thought it would be easiest for our workflow and the network.

So you've minimized FireWire in your workflow?

We still store our high-res QuickTimes on FireWire drives. Our company at any given time

could be doing between eight to 14 television series so there's no way to cut all of our shows in full-res. Instead, we take the QuickTime files, recompress them in Final Cut Pro, but make our own proxies that retain the same timecode as the original hires QuickTimes. When we have time to rebuild the show, we'll go back to those FireWire drives and begin a media recovery/relinking process. We've had a pretty successful run on this series, but we've had two FireWires in the past that, months after archiving to them, failed to light up. Then, luckily, we had the LTO tape that we stored that data on the Cache-A with, in order to rebuild that show. The Prime-Cache has saved us many a time.

Why was the LTO tape component appealing to you?

LTO tape is pretty much universal to the networks. On top of that, it was a very easy installation. Also, the ability to connect it to all of our clients and archive from anywhere in our building was a huge selling point because we do so many shows at our company. There's no way we can have one device connected to one computer. That computer will never always be available.

How many people use the system on a daily basis?

Four assistant editors, who all back up to the unit.

Will you use the Prime-Cache for HD on any future projects?

Actually, yes. Our long-running hit show on the Food Network, *Ace of Cakes*, is going HD for the first time in spring 2010. We want to shoot it on the Sony EX3 or Panasonic P2, the reason being that our camera operators have to work in a very stealthy manner inside the bakery. If they have big cameras with ENG units on their shoulders, they may tend to bump into



the people that work there. So using file-based cards is the way to go for that series.

I'm in the process of developing a new workflow for *Cakes* because that show in particular shoots a tremendous amount of footage — they need to shoot lots and lots of time-lapse footage of the cakes being prepared. It probably shoots three or four times what other series do. So I don't think it would work to unload FireWires in the field, because of the huge capacity, probably about 180 TB of FireWire. For them to properly store all this footage, I'd like to send out the newer Cache-A Pro-Cache archive appliance*, which is larger, into the field and have two digital technicians archiving day and night onto LTO tape. Shipping FireWire back and forth across the country is worrisome—are they going to work when they arrive? I have the security that LTO tapes will survive the journey. ■

**At press time, Cache-A had just debuted the Prime-Cache Pro appliance at NAB 2010.*

Jessica Clegg, formerly an editor at Studiodaily.com, is the managing editor of CableFAX Magazine.

Creating a Shared Tiered Storage Archive

How RLTV Kicked the Videotape Habit

By Bryant Frazer

For James Peebles, VP of engineering at Retirement Living TV (RLTV), building a new automated, tier-based archiving system from the ground up meant an opportunity to put the finishing touches on concepts that he started developing in 2000, when he worked for Oxygen Cable. As the emergence of workgroup editing created shared storage pools that were exploding in size, Peebles realized that existing data-management systems weren't up to the task of handling 21st century workflows. "The easy part is hooking up your editing machines and getting them to work together," he says. "The hard part is managing the data you're acquiring—and managing it in such a way that you can use it later and organize it sensibly."

The problem, in other words, isn't just storing footage on ever-larger volumes. You need a strategy that gives your storage system information about itself, and instructions on how to handle all those bits and bytes in a way that computes in terms of bandwidth requirements for backups, flexibility when it comes to ingest and playout, and, last but certainly not least, cost-efficiency. There are lots of ways to build that kind of system, and each one has its ups and downs. But Peeble and his team at RLTV dedicated themselves to building a system that would not just keep up with technological change, but with the cultural change that comes with the arrival of totally tapeless workflows.

"We have a truly file-based workflow, and we're one of the few networks we know of that has completely eliminated videotape from the process," Peebles says. "A lot of networks are trying to reuse and repurpose legacy systems, with a production style that is built on videotape with producers and

editors programmed to work in that mode. We've shaken all that off, retraining our editors and producers to work in an all-digital environment."

Peebles knew he wanted to move RLTV away from the industry's videotape legacy, which meant archival capabilities and disaster recovery would be a key issue. The first question had to do with what general components would make up the system—specifically, did it make more sense to go with an all-disk-based system, or to try to integrate tape backups with spinning disks? RLTV took a rigorous look at pure disk-based solutions, noting that companies like Isilon offer well-engineered, expandable systems that allow storage to grow organically as a facility's needs increase.

But Peebles found that the overhead costs of keeping a large spinning-disk system running would be substantial, and wondered about the implementation of a backup system for disaster recovery.

"As you build up more banks of disks, you're having multiple devices running simultaneously, drawing tremendous amounts of power and creating tremendous amounts of heat. That has a lot of costs," he explains. "And backing up the system puts you on the horns of a dilemma: if you're going to stay all spinning disk, you have to replicate that system somewhere else, with a data pipe copying data from one physical location to another, so that you have true redundancy. That's a very expensive solution that requires a lot of horsepower, a lot of electricity, and a lot of cooling. It's all moving parts, all the time."

Peebles decided instead to lean on LTO tape backups, mainly because the media only has to be loaded and spun up to speed when it's actually being used by the system. And, if you can instruct your storage

solution to write everything that goes to data tape onto a second physical tape unit, you can physically transport those tapes to a secure off-site location.

RLTV eventually decided to put a small amount of disk-based storage at the front end of the system, with a large data-tape vault on the back end. It wasn't hard to put the system together, and Peebles approached vendors he had worked with in the past. Apple provided its Xserve and Final Cut Server technology, Quantum provided its StorNext data-management software and the Scalar i2000 tape library, QLogic provided fiber-switching infrastructure and Cisco Ethernet switching, and Promise Technology became the go-to provider as RLTV needed to expand its disk-based storage capacity.

Rock Around the Clock

"The original design intent allowed for two basic things," says Andrew Richards, RLTV's director of engineering technology, who worked closely with Peebles to build the system. "One was editing, if necessary, around the clock on multiple workstations, all sharing storage. And the other was to maintain file-based output rather than creating tape masters of the finished content. In order to pull that off, we needed to have a shared-file system that could cycle its data to a nearline platform without impacting the performance of the system to allow editing to go on." That meant traditional backup techniques, which use up a chunk of bandwidth just to transfer copies of data to multiple tape drives, were out of the question. Instead, the StorNext software looks at the disk-based storage, file by file, as new data is created and replicates the new pieces of data to tape. Essentially, the backup is constantly running — but never using up enough bandwidth to impact the editorial process in multiple rooms.

"It's an intelligent system," Peebles says, "unlike traditional backup systems where an operator has to get involved in physically managing your backups and incrementals. We tell it, 'when you make a copy, make two copies. We take one out and put it in a cardboard box in a secure location, and that's the vault.'" Peebles had tried to make similar systems work before, but he says the technology just couldn't keep up with itself. "When we moved to our first SAN systems at Oxygen Cable, I tried to do backup systems that would work, but the rate at which data was changing was so high that there

"We have a truly file-based workflow, and we're one of the few networks we know of that has completely eliminated videotape from the process,"
Peebles says.

wasn't a tape robot in the world that could keep up with a backup."

Richards explains further. "The metadata of the file system that holds the data drives the policies of moving data to tape. That happens at such a low level that it's not going to overwhelm the bandwidth of the system the way old-style overnight backups would. It's just quietly trickling away, keeping the data safe. The policy we have set up will put the data on tape within 15 minutes of its being introduced to the system.

Final Cut Server in the Mix

Final Cut Server came into the picture, two years ago, as an exceptionally functional digital asset management solution. StorNext and FCS don't exactly talk to each other, Richards says, but they work exceptionally well together. "Part of our build-out for Final Cut Server included the desire to completely blow up our file system and start with a fresh arrangement to complement what FCS could automate for us," he explains. After the release of Final Cut Server in 2007, RLTV shrank its existing file system to fit on some newly purchased disks, then repurposed the original Apple storage system to build a larger file system from the ground up. "That brought us into a better logical arrangement, with FCS managing the creation of directories for us against projects the teams were starting. It also manages the movement of finished content on and off of nearline storage, based on user input. Before FCS we had to manually interact with the file system just to get data back off of tape for people. But

now, they're just a few mouseclicks from recovering content from the nearline tape archives."

When FCS moves a file to a given location, StorNext takes the ball, moving the file to tape and deleting it from the disk. If FCS comes looking for that file again, it just has to wait for it to be retrieved by the Scalar i2000 library and then put it back online. (One of the best features of the i2000, according to Richards, is its highly predictable cost. "We chose that platform because it could expand into the petabytes if necessary," he says. "Every time you buy another 100 slots, regardless of whether or not a new cabinet is required to accommodate them, it's always the same price. That was appealing — we always know how much it will cost to expand, regardless of how full the hardware is.")

And Final Cut Server turned out to offer extra capabilities that the team at RLTV wasn't even aware of when it purchased the system. "Because of the way it hooks into Final Cut systems like Compressor and Qmaster, Final Cut Server has become part of our distribution engine," Peeble explains. That means FCS can be used to transcode video for use on the Internet, mobile phones, and elsewhere. "We run scripts from FCS to create and distribute media by FTP anywhere in the world. It became an opportunity to fully automate the process — when we are responsible for acquisition, we don't even go to tape. We build spinning-disk-based acquisition systems that take us straight into the system. And if we take acquisition from outside the system, we have enough decks to ingest it once, and it never touches videotape again. It's all file-based from then on, for ever after."

Set It and Forget It?

It took a lot of work in the planning stages, but now that the system is in place, it requires remarkably little care. "It took us a while to get it running perfectly, but now that it's running it takes care of itself," Peebles says. "We even had a catastrophic event — we were in a temporary location where the air-conditioning system failed over a weekend and cooked the whole system to 120 degrees. It cor-

rupted the system, which we had to blow up and rebuild. So we actually proved that disaster recovery works — the hard way."

"This was a file system of 25 TB," Richards recalls. There was downtime to actually recover the file system, but the data was intact."

The main production file system at RLTV today is 35 TB, with an additional 10 TB set aside for swap files and proxy storage. The Scalar i2000 is licensed to hold up to 85 TB, and the tape vault holds about 200 TB worth of data.

One of the reasons this is all so affordable is because RLTV is standardized around DV25, which is still standard definition, but can be easily run through a Teranex converter that will make it look exceptionally good at HD resolutions. "We're poised to make a transition to ProRes or DV100, but we're still DV25 until there's a demand from our distribution side to convert," Peebles says, noting that the system is integrated with playout. "Master Control plugs into our storage system. We move our finished content into our playout system, segment it, and play it to air without any tape intervention at all. Frankly, we're capable of moving it all to the Internet if we want. Because we're fairly new as a company, we have been able to build this kind of system from the ground up."

As companies grapple with the HD transition, Richards suggests they should consider making a break with videotape-based tradition. "If they're going HD, they can go file-based at the same time and shrug off a lot of old habits and techniques," he says. "Culturally, there can be a lot to overcome. There's industrial inertia from people who have been doing this for decades. But the cost advantages are so significant that it's impossible to ignore going to a file-based technique like this." ■

Bryant Frazer is the editor-in-chief of Film & Video and co-editor of StudioDaily.com, Studio Group publications dedicated to production and post-production for film, television and online. His writing has appeared in Studio/monthly, AV Video Multimedia Producer, Tape Disc Business, and DVD Report.

Restoring and Archiving Film for DVD, Web, Mobile and Beyond

Canada's Historic Visual Vaults Prepare for the Digital Future

By **Beth Marchant**

For the past 71 years, the National Film Board of Canada has been preserving the social and cultural history of Canada on film and video, largely in the form of documentaries, animations and alternative drama. The government-run agency—created by an act of Parliament and alternately known as the NFB and in bilingual Montreal, Quebec, where it is headquartered, as the Office National du Film du Canada, or ONF—is Canada's public film producer and distributor. It operates just like a production company, staffing an extensive in-house post pipeline in top-tier scanning, editing and finishing suites across its national offices. Of the 13,000 videos and films to its credit, more than 100 have taken film's highest honors, including 12 Oscars and 90 Genie Awards from the Academy of Canadian Cinema. The NFB has amassed, in fact, more than 5,000 awards since opening its doors in 1939.

But that's just part of who the NFB is and what it does. Its larger mission, says general manager and COO/CFO Luisa Frate—who reports to the NFB's chairman of the board, the Government Film Commissioner—is to store, catalogue, restore and dynamically share all these national treasures, plus future works that are rediscovered and come into its hands, with Canadians and the world. In 2009, the NFB began making its collection available free online in multiple formats, from streaming videos on its Web site via its "NFB Screening Room" (<http://nfb.ca/>) to streaming content through iPhone and iPad apps it developed with software partners. Though not yet tapping into the NFB's entire collection—some 6,000 hours of footage, at last count—each app streams hundreds of NFB



The NFB is in the process of digitizing all of its original film footage into multi-resolution mirror files for archiving, restoration and distribution.

films and videos to those devices. The NFB's online store also continues to encode and sell DVDs on demand for the Canadian, European and U.S. market in both the PAL and NTSC formats. Most recently, the NFB added HD and 3D films to the collection, for a total of more than 1400 viewable works online.

All of this would be impossible to manage, says Frate, without a clear digital strategic plan and a deep and nimble archive. "In the past, we simply



Luisa Frate, the NFB's General Manager

backed our content up, but we never had any kind of full-fledged archival system, where we could restore the information and retrieve it at will," she says. "A while back, we had a grant, called "The Memory Fund," where we tried to put several of our films online. But that was just a small window of what we are offering online now."

Since the latest commissioner, Tom Perlmutter, took office in 2007, says Frate, "we have really focused on making a complete digital shift. Internally, we really oriented everything, including our budgets, to prepare ourselves for moving all of our productions online in some digital form." The goal, which is already well underway, is to restore and archive within the next 18 – 24 months every original work digitally at 2K, 4K and 6K resolutions and make it available in any format on demand. Approximately 80 percent of the NFB's collection is on film, with the remaining 20 percent recorded on the leading various videotape formats of the past several decades.

Upgrading on Budget

It had always been a challenge for the NFB, which reports to the Canadian Parliament through the Minister of Canadian Heritage, to keep up with accelerating technology changes while navigating the tightly controlled and scrutinized budgets that come with being a government agency. But, says Frate, instead of delaying the NFB's emergence

into the digital age, those limitations have actually helped the film board accelerate both its digital plan and its daily workflows. "We became as efficient as possible early on so we could purchase the equipment we needed," she says. It's also made it much easier, she says, to be equally open and transparent about the digital transition with the public.

The NFB also has a robust research & development department, comparable to those in many public technology companies, that helps accelerate its restoration, archiving and deployment efforts in house or with the help of technology partners. Last year, the NFB began working with Atempo, the Silicon Valley- and Paris-based archiving company, to come up with a way to link its custom asset manager, built –in-house ten years ago, to its wider work of restoring and archiving its collection.

"We were looking for a partner who could help us catalogue our vast resources and create a search engine that was robust enough for all our different users to find what they were looking for in various languages," explains Julie Dutrisac, head of R&D. The first step, she says, was to purchase and install an Atempo Digital Archive. "But we also needed help creating and implementing a structure that was completely adapted to our needs," she says. "We've been digitizing for some time, but not always in the best format and resolution. Thanks to the Memory Fund, we've been digitizing from videotape to expand the access to all users. More recently, for about one year, we have been upresing and restoring our films from 2K to 4K, and all the way up to finalizing a restored digital master."

The complete process, she says, has involved installing the equipment, capturing the content, optimizing the amount of data that is captured, and making sure, above all, that everything is tagged with the proper metadata. "For film, we telecine and digitize every single element separately and create a DSM, or Digital Source Master, which is sort of a clone of what we have on film. For example, when digitizing 100 films, we generate 10 million metadata entries and approximately 30 TB of content. If you tagged all of those as separate files, that would be approximately 2.4 million different files. This is all archived, image for image, so we have a lot of data going directly from our storage to the Atempo system just to create a DSM."

Dutrisac and her team worked with the Atempo R&D team to write an XML-based plug-in to simplify the process of moving every last bit of that

data, as various versions of a film, throughout their existing pipeline and up into staff member's workstations through precise and easy metadata searches. "The plug-in handles all of our hundreds of millions of metadata transfers across our normal restoration and digitization workflow and asset management system, and then stores, indexes and manages them within the Atempo Digital Archive." The plug-in supports a range of media files, including DPX.

Mirror Versions and Metadata

After creating a Digital Source Master and pumping that data directly into Atempo with the custom plug-in, the NFB team moves through a traditional restoration workflow using off-the-shelf tools, from color grading to synchronization, and simultaneously considers the multiple versions of the film or animation they need to create. All of it, says Dutrisac, is integrated with the Atempo Digital Archive.

"We've been doing versioning for a long time now," she says. "Most of our films are already close-captioned in both English and French, since by law we have had to do that since the 1980s. Some films have already been translated into many more languages. Now we are dealing with programming in the video descriptions into the closed-captioning, targeted at specific age groups, from young children to adults. We need to figure out all the elements that are in the vault so we can generate all these versions completely. We also sometimes create different durations of the film," to be able to deliver shorter streamed versions online or on handheld devices, she says.

Only when the films are restored, versioned and color graded does the NFB create the final digital master, which is also stored with the other versions within the Atempo system. Says Frate, "Every piece of information is integrated and linked to our asset management system in order to ensure that we have access, at any time, to the actual film or video itself. Within our three-step digitization process, where the first step is just to digitize the film, we've got a mirror of the original. If there are scratches and snow on that original, they are recorded on the Digital Source Master. The restored film, with the color enhanced and defects removed, lives on the same system."

How does the NFB save server space with multiple copies coexisting within the same archive? "The deliverables themselves don't exist anywhere on the Atempo," says Frate. "Our asset manager, in the past, was mostly focused on finished goods, like DVDs. Now we're shifting to a virtual product and video-on-demand, encoding films only when we need to for a DVD order, or for our iPhone partner to deliver as part of the app, or if we are running a promotion on NFB.ca on a particular topic. That alone saves us incredible amounts of time and storage capacity."



Julie Dutrisac, head of the NFB's R&D team

In addition to the uncompressed Digital Source Masters, the NFB also archives each film's restored Digital Master, or DM, for theatrical screenings, as well as a mirror version encoded as a mezzanine file, an unassembled, partially compressed file format that can be later transcoded on-demand for DVD, SD and HD broadcast, or for any number of online or mobile streaming outputs. "That is why our asset management and archival system has become so critical to our workflow," says Dutrisac. "We are not just keeping track of every piece of original content but also all of the mezzanine file recipes for each version of our titles, so we can reassemble them on demand at a moment's notice."

It's a complex and often daunting process for any media facility, let alone one NFB's size. "We've been working hard recently to automate these workflows so we can process as many films as possible in a short period of time," she adds. "So for us, given our staff, the bigger challenge is time. We're handling this by programming special applications



Restoring a film inside one of the NFB's DI suites

to be able to be more efficient. But as Luisa mentioned, because of the way we encode on demand, we're able to manage all of our content with a relatively small amount of storage, even though resolutions are always getting higher and files continue to get more and more metadata rich."

The Tapeless Future

Frate says although the NFB is not yet completely tapeless, that is the near-term goal. "Julie and her team are working very hard to get our workflow tapeless from end-to-end within the next one or two years." The NFB is also currently working to move the Digital Source Master process completely into the Atempo Digital Archive workflow.

With an eye toward future-proofing the NFB's workflow, adds Dutrisac, "we're always evaluating new applications and software and hardware. But right now, we are getting exactly what we need from the Atempo solution." [EDITOR'S NOTE: As a Canadian government agency, the NFB is prohibited by law from publicly naming any branded product or service it purchases for use. Business and development partners, such as Atempo, are not subject to this rule.]

"The flexibility of the Atempo system got our attention, and by partnering with Atempo's R&D, we could seamlessly connect the fuller archival workflow to our existing asset management system," Frate says. "And we're not the only beneficiaries of this plug-in, either. Now Atempo can use that

plug-in we co-developed with its other customers."

Another huge learning curve, Frate says, has been the way NFB staff members work together within this evolving process. With NFB offices across Canada, the next challenge is to pull the entire organization into the integrated archive across the network. "We work with these centers at a distance. Our next phase of the process is to be able to archive our Vancouver office with our Montreal office. We are not there yet, but we are certainly

starting to put the preliminary steps in place."

How far Canada's National Film Board has come digitally, given its limitations, speaks volumes about its small yet motivated and talented staff. "The NFB isn't that big and we have a relatively small budget, yet we've been able to think about, afford and use existing technology by modifying it so that it makes sense for us, automating things that will save us time and resources," says Frate. "We're not a Warner Bros. A lot of organizations already do this kind of customization, but I think we often go one step further in terms of programming and automation so we can make great technology work within our budgets." ■

Beth Marchant is the founding editor-in-chief of Studio/monthly and the co-editor of StudioDaily.com, a collection of Web sites, including Studiomonthly.com and filmandvideomagazine.com, dedicated to production and post-production technology. She has spoken on a range of digital content creation topics at the National Association of Broadcasters convention, the Film Finance Forum and at the Writers Guild of America. Her articles and photographs have appeared in Yankee, Harvard Magazine, on the syndicated Gannett News Service and in numerous Studio Group publications in print and online. She also reviews consumer technology for the site Techlicious.

For all its advantages, the transition to digital workflows has also come with its challenges: bloated storage, unprotected assets and proprietary formats.

3

New Directions

Ready for Prime Time

The Digital Archive's New Leading Role in Media Workflows

By Janet LaFleur

Across all segments of the media and entertainment industry, the past decade has seen a fundamental transformation, where virtually the entire industry moved from tape or film to digital media. This change required investing in digital equipment and processes: new cameras and new edit workstations, additional servers, storage and software, and hiring or training staff for these new workflows. While the transition to digital is complete for most, the work is far from over. The transition to HD and other higher definition formats is the next wave that has already hit many facilities. Despite the advantages of digital workflows, there are significant challenges.

Expanding Storage Volumes

With each hour of HD footage requiring up to six times more storage capacity than standard definition, the transition from SD to HD is causing content creators to rethink their existing storage architecture. An inadequate storage infrastructure can leave editors scrambling to free up space before they can start editing a new project—or even turn down higher-definition projects for lack of storage capacity. Sadly, in many cases the capacity exists but is tied up in “storage waste,” content that is duplicated within the system as editors swap files and neglect to clean up.

Unprotected Content

With new high-definition digital cameras using more expensive, reusable hard drives or flash as capture media, the natural back up that the original film or videotape provided no longer exists. Without a proper data protection strategy, raw footage is left at risk. A disk crash or accidental

deletion can mean a permanent loss of irreplaceable footage.

Proprietary Formats

As digital media workflow technologies emerged, a host of proprietary formats and devices were established with them. Proprietary technologies can hinder media companies from rolling out new technologies and new processes or, worse, can create vendor lock-in.

For all its advantages, the transition to digital workflows has also come with its challenges: bloated storage, unprotected assets and proprietary formats. By giving archiving a new, central role in digital media workflows, organizations can address many of these challenges. This chapter outlines six strategies that broadcast and post-production shops can use to protect digital assets, streamline workflows and optimize storage.

Rethinking the Role of the Archive

For all these data management challenges, there is a solution—the digital archive. But to solve these challenges, the definition of the archive needs to be expanded. Traditionally, archiving was the last step of the creative process. Just before broadcast or distribution, the content would be archived onto a master tape or film would be stored in a media vault. These vaults were intended as repositories that could be manually searched to retrieve content for later re-use. The harsh reality is that these vaults have been more burial ground than archive, prompting one pundit to label them as “a place old clips are sent to die through lack of use, an offline video cemetery.”¹ A well-architected digital archive can be much more than a final resting place. Instead of a final step in the edit process, the archive should be an integral part of the digital media workflow. In a March 2008 seminar, Samuel

Reisner of ESPN Sports underscored this concept, stating, “Any piece of media must be available on any production platform whenever it is needed.” From ingest to edit, to playout or distribution, offering archive access to all steps in the workflow achieves a range of objectives.

Content Re-Use

An online digital archive, especially one with rich search and retrieval capabilities, keeps previously created content at producers’ or editors’ fingertips, making re-use quick and easy. Managing storage growth data migration or hierarchical storage management (HSM) features allow administrators to align the storage performance (and cost) with the access needs and relative value of the content. Making content available doesn’t necessarily mean it needs to be available on high-performance storage. Moving less frequently used content to tape or lower-performance disk can significantly reduce infrastructure costs.

Protecting Assets

The digital archive protects and preserves content throughout its lifespan—from backing up otherwise unprotected raw footage to becoming the master vault for historical preservation projects.

Six Strategies for Centering the Archive in your Workflow

With the new, expanded role of the archive, what strategies should organizations take to maximize its potential?

Strategy 1: Preserve Raw Footage in the Archive

With creative content, you often don’t know what to keep and what to leave on the cutting room floor. Even the most ordinary clip can become priceless overnight. When the scandal over Bill Clinton and Monica Lewinsky broke, news organizations scoured their archives, racing for footage or still photos of the two together. The lucky winner? CNN, who dug up two-year-old footage of the couple in their infamous embrace. An obscure clip of an uneventful crowd scene became a hot news item that influenced the world’s perceptions of the controversy. Archiving raw footage is the best way to ensure it will always be available. If you choose this strategy, you should be able to:

- › Automatically archive new footage as soon as it enters the workflow, using precise criteria

such as file location, file type, creation date and last access date.

- › Archive to inexpensive media like digital tape in a non-proprietary format like tar.
- › Attach metadata to archived files for future search and retrieval.
- › Retain the files on high performance storage for editors to use.
- › Automatically delete the files from production storage after a configurable period of time.

Strategy 2: Empower Users to Archive, Search and Retrieve Directly

Eliminate the middleman and give content creators direct access to the archive. Once the archivist or librarian sets up the archive structures and policies, editors should be able to archive their projects as soon as they’re completed. With direct retrieval, archived content can be immediately available to editors or producers for reuse in new projects. Choosing a solution with an end-user interface for archive and retrieval will allow your archive to scale without overloading your archivist. Some additional features to consider:

- › A drag-and-drop graphical interface that displays both the user’s file system and archive’s hierarchical structure and allow the user to drag and drop files into or from the archive’s folders.
- › Search capabilities that allow search based on a broad range of criteria, and allow users to directly retrieve projects to any location, regardless of where the project originated.

Strategy 3: Avoid Proprietary Workflow Components

For many facilities, the transition to digital workflows brought along some unwelcome guests: closed, proprietary solutions. From closed file formats to hardware appliances to full turnkey systems, proprietary workflow components restrict a facility’s growth, innovation and evolution toward new technologies. To maintain your environment’s ability to choose best-of-breed workflow solutions, consider a solution that:

- › Provides support for archiving in standard, open file format like tar.
- › Is integrated with open asset management solutions like Final Cut Server.

- › Offers an open API for ease of integration with other management tools.
- › Is sold as software, rather than as an appliance, which limits hardware choice, scalability and upgradability.
- › Supports a variety of archive media and platforms, including disk, tape and even cloud storage.

Strategy 4: Use Automatic Archiving to Reduce Clutter in Work Areas

Despite users' best intentions, work areas on edit workstations and shared storage areas tend to grow out of control. Like an office refrigerator that everyone uses, but no one cleans out, collaborative work areas can get very crowded with leftover, forgotten files. Automated tasks that search for stale data and migrate it to an off-line archive can free up valuable production storage space. To reduce clutter and reclaim wasted storage, look for an archive solution that:

- › Uses site-specific criteria like last access date and file type to automatically archive outdated files.
- › Once files are archived, removes files from the work area.
- › Allows users to retrieve from the archive any files that should not have been removed.
- › Uses de-duplication to optimize available storage space.
- › Does not rely on stubbing as the only retrieval method. Stub files may reduce the storage volume, but they don't truly reduce clutter since the file names still appear in the directory.

Strategy 5: Set Up Multi-Tier Archives

As your archive expands to protect and preserve all your digital assets, the reality is that only a subset may be retrieved from the archive. The challenge is predicting which assets may be required and being prepared when they're requested. An archive that spans multiple tiers of storage can reserve more expensive near-line storage for content that requires higher performance, and uses a less expensive, offline "deep" storage for long-term preservation or lower priority content. To enable you to match your content's access requirements to the storage performance, your archive solution should:

- › Simultaneously archive to both near-line and deep archives either manually, through

an automated task or from an asset management system.

- › On retrieval, grab the content from the fastest storage available.
- › Automatically remove an item from near-line storage according to site-specific criteria, such as after six months of disuse.
- › Retrieve the asset from the deep archive after the file has been removed from near-line storage.

Strategy 6: Plan for Large-Scale Preservation

From call-letter stations to media giants to government agencies, digital preservation projects are being launched to protect historical films, cultural treasures, and other valuable assets while they're still accessible. Their plan: open up their extensive tape and film vaults for content for re-use, re-distribution or as a public resource. Digital preservation projects protect content from a host of hazards. Film and videotape degrade over time, playback devices become obsolete and break and formats evolve, leaving content unreadable even if the media is intact. To prepare you for large-scale preservation projects, your archive solution must:

- › Support high-performance configurations to allow massive numbers of files to be archived daily.
- › Scale the archive itself to hundreds or thousands of terabytes.
- › Use non-proprietary formats and technologies that support a broad range of media and device types.

Janet LaFleur is the director of field marketing at Atempo in Palo Alto, CA. The software-based Atempo Digital Archive uses open formats (tar format on tape and regular files on disk) to archive to a broad range of devices, including tape, disk and cloud storage. It is integrated directly with Final Cut Server, letting you archive from within the Final Cut Server interface. For more details about Atempo Digital Archive's scalable, flexible and open approaches to managing your content, visit www.atempo.com.

Sources: ¹ *The Register*, "Apple eats video editing jobs," http://www.theregister.co.uk/2008/12/01/video_editing_online_archives/. ² Samuel Reisner, ESPN, "The Sixth Easy Piece" Seminar presented at Netherlands Institute for Sound and Vision, March 6-7, 2008

4

Resources

Resources

A directory of primary manufacturers and companies offering archiving-related products and service

Linear Tape-Open (LTO) Drive and Drive Library Manufacturers

Dell PowerVault Series

www.dell.com/

HP StorageWorks Ultrium Series

www.hp.com/

IBM Ultrium Series

www.ibm.com

Oracle Sun StorageTek Series

(formerly Sun Microsystems LTO drives)

www.oracle.com

Quantum Ultrium Series

www.quantum.com

Seagate Viper 200 Series

www.seagate.com

Tandberg Data

www.tandbergdata.com/

LTO Tape Media Manufacturers

Fujifilm

www.fujifilm.com/products/storage/lineup/ltoultrium/

HP

www.compaq.com

IBM

www.ibm.com

Imation

www.imation.com

Maxell

www.maxell-usa.com

Oracle (Sun StorageTek LTO Tape)

www.oracle.com

Overland Storage

www.overlandstorage.com

Qualstar

www.qualstar.com/

Quantum

www.quantum.com

Sony

www.sony.com/professional

Tandberg Data

www.tandbergdata.com

TDK

www.tdk-media.com

Verbatim

www.verbatim.com

Additional LTO Format Resources

The LTO Program

www.ultrium.com/

Consortium formed in 1997 by HP, IBM, and Seagate (now with Quantum in its place) offers technology overviews, white papers, news and webinars about the Ultrium LTO format.

Blu-ray Disc (BD) Media Drive and Drive Library Manufacturers

Buffalo

www.buffalotech.com

DelkinDevices

www.delkin.com

Denon

www.usa.denon.com

Hitachi/LG

www.lge.com

JVC

www.pro.jvc.com

La Cie

www.lacie.com

LG Electronics

www.lge.com

Lite-On

www.liteonit.com

Microboards

www.microboards.com

Panasonic

www.panasonic.com

Pioneer

www.pioneer.com

Primera

www.primera.com

Sony

www.sonystyle.com

Tascam

www.tascam.com

Yamaha

www.yamaha.com

Blu-ray Disc (BD) Media Manufacturers

Maxell

www.maxell-usa.com/

Memorex

www.memorex.com/

Optical Quantum

www.opticalquantum.com/

Panasonic

www.panasonic.net/avc/blu-ray_disc/

RIDATA

www.ritekusa.com/

Sony

www.sonystyle.com/

TDK

www.tdk-media.com

Verbatim

www.verbatim.com

M-DISC Drive and Media Manufacturers

Millenniata, Inc.

Makers of the M-WRITER and M-DISC, a DVD-readable disc with a rock-like coating that is purported to last for 1,000 years. Current capacity on one disc is 4.7 GB. The M-WRITER drive is a portable external drive with an industry-standard USB 2.0 Interface. www.millenniata.com/

Primera

Bravo Archive-Series Disc Publishers. The recently released Bravo SE and BravoPro Xi2 features Millenniata's M-DISC technology. www.primera.com/millenniata

LTO, BD and M-DISC Drive, Library, Cartridge and Case Resellers

CDW.com

Tapeandmedia.com

Tape4backup.com

Overstock.com

Tapelibrary.com

Supermediastore.com

www.tapebackupoutlet.com

www.datatechstore.com

bhphotovideo.com (B+H Photo)

tapereources.com

macmall.com

rmeinc.com

SCSI Storage Controller Card Manufacturers

Adaptec

www.adaptech.com

ATTO

www.attotech.com

Quantum

www.quantum.com

StarTech

www.startech.com

LTO and Disc Back-Up Software

AASync FTP Sync

SFTP synchronization and file backup tool for Mac OS X and Windows.

www.aasync.com/

Atempo Time Navigator

Web-based tool for centrally managing backup and recovery across a production workflow and archive. One backup stream enables multiple copies of a file both locally and remotely.

www.atempo.com

CA Software ARCserve Backup

Backup software with built-in data deduplication.

arcserv.com/us/

Barracudaware Yosemite Server Backup

Backs up Windows, Linux, Netware to disk and a variety of tape-loading devices.

www.barracudaware.com/

CommVault Galaxy and Simpana

Includes modules for backup & recovery, archive, replication, resource management and search.

www.commvault.com

HP Data Protector Express

Backup software aimed at small facilities without a specialized IT staff in house.

h18000.www1.hp.com/products/storage/software/datapexp/index.html

Imagine Products ShotPut Pro and ShotPut RED

Offloading and management software for P2, Sony SxS, AVCHD, JVC ProHD and RED ONE that

makes secure, auto-named copies of a card or hard disk's contents in up to three locations.

www.imagineproducts.com/

R3D Data Manager

From on set, to editorial to archiving, R3D Data Manager verifies that RED footage is correctly copied and rendered, putting copies in up to four locations. Can re-check files against any source at any time, even if the original has been deleted, and render R3Ds to QuickTime.

www.r3ddata.com/

Symantec Veritas Backup Exec

Features integrated deduplication and centralized 3-tier setup, reporting and patch management.

www.symantec.com/business/backup-exec-for-windows-servers.

Tolis Group BRU Server

This backup and restore software features live database back-up on server and client systems.

www.tolisgroup.com/

LTO-, Blu-ray- (BD) and Tiered Storage-Attached Archive Appliances and Tape Libraries

1 Beyond Wrangler Pro

Portable tapeless ingest/edit station that features optional built-in DLT or LTO-3 and LTO-4 drives for archiving.

www.1beyond.com/index.asp

Alteran Technologies ViTaDi

Suite of tools that captures, transcodes, and delivers video in multiple video/audio file formats to media servers, digital archives and asset management systems.

www.alterantech.com

Cache-A Prime-Cache and Pro-Cache

Network-attached archive appliances that provide both source masters for digital acquisition and long-term storage for project archives using LTO-4 and LTO-5 tape.

www.cache-a.com/

Front Porch Digital SAMMA Solo

A semi-automated video-tape migration appliance for processing high-volume encoding and metadata.

www.fpdigital.com

Focus Enhancements ProxSys MA Series

Integrated ProxSys media transfer tool imports native metadata and archives automatically to dual BD drives.

www.focusinfo.com

Integrated Media Technologies Archive Systems

Uses LTO and Titanium drives and hierarchical storage management systems integrated with a broad range of disk-based technologies to deliver storage and backup for technology applications in video, film production, and post production.

www.imtglobalinc.com/media-entertainment/storage/archive-systems/

Media Distributors Archive Station

Portable turnkey, Mac-based real-time hard-drive backup that features an LTO-4/5 drive or a Blu-ray burner for archiving and Constellation VCM Asset Management and Final Cut Studio 2.0 Software

www.mediadistributors.com/access/index.php?/site/rentals_thearchivestation/

Nexto Video Storage Pro

Portable backup device for Sony SxS and Panasonic P2 cards.

www.nextodiusa.com

Qualstar BQ Series Video Archive Systems

BQ Series VAS includes archive servers configured for popular archive software packages, large disk cache for fast access to current content and a robotic tape library.

www.qualstar.com/146086.html

Rorke Data Digital Content Archive (DCA)

DCA for Video is based on Rorke's Galaxy RAID 6 storage server, digital tape libraries and its HSM asset data mover software application. DCA for Film is built to manage and archive 2K, 4K and 3D stereoscopic files and supports disk and tape for short and long-term retention.

www.rorke.com

Sony XDCAM Cart and PetaSite

IT-based network archive (XDCAM Cart) and scalable tape library (PetaSite) for managing and archiving larger amounts of footage.

pro.sony.com/bbsc/ssr/app-archiving/

Spectra Logic Spectra T-Finity

Uses multiple, redundant robots, scaling to more than 45 petabytes in a single library and to more than 180 PB in a single library complex; aimed at large enterprise IT, federal, high performance computing (HPC) and media/entertainment.

www.spectrallogic.com

Tandberg Data

Backup, deduplication and disk- and tape-based (LTO) libraries and storage appliances.

www.tandbergdata.com/us/

Tolis Group bruAPP Backup Appliance

New appliance that incorporates BRU Server backup software and LTO drives in 1U and 3U rack-mount configurations. Supports Unix, Mac OS X and Windows and claims to have the lowest cost per seat of any true client/server-based backup solution.

www.tolisgroup.com/

Media Asset Management (MAM) and Content Storage Management (CSM) Software and Servers

Accordent Media Management System

The Accordent Media Management System is centralized online infrastructure for ingesting, organizing, searching, publishing and securing multi-level access to Accordent presentations and other multimedia communications, including archived Web- and video-conferences. AMMS Express launched in March 2010.

www.accordent.com

Active Circle Archive Edition

Archive Edition is based on Active Video Archive, which manages archives on disk and tape while improving ease of access. All Active Circle archiving solutions are based on standard IT hardware to keep costs down and are scalable and easy to integrate into a digital workflow.

www.active-circle.com

Apple Final Cut Server

Asset sharing and managing tool for Final Cut Pro-based production and post environments. Can be used with software-based archiving systems such as Atempo's Digital Archive or Quantum's StorNext. www.apple.com/finalcutservers

ARCHIWARE PresStore

ARCHIWARE's data management solution PresSTORE lets users synchronize, backup and archive data in heterogeneous environments. www.archiware.com

Avid Interplay/Integrated Media Enterprise

Customizable DAM for managing workflows for networked and remote users; open architecture supports all Avid and non-Avid file types and applications. New in 2010: Integrated Media Enterprise framework, which controls Interplay Production and Archive through a workflow manager, is designed to work across wide area networks. Avid acquired in March 2010 Blue Order, a product that transforms Avid's production asset management technology into a full media asset management product. www.avid.com/products/Interplay/index.asp

Avid Interplay Archive

Networked service that lets Interplay users search, restore and browse archived content and restore to the online library for immediate work. Driven by SGL FlashNet system. www.avid.com/products/Interplay-Archive/

Avid Unity Media Network

Designed for storing, accessing, and sharing high-resolution creative media in collaborative, real-time workgroup environments. www.avid.com/products/Unity-Media-Network/index.asp

Bitcentral Oasis

Browser-based media asset management system aimed at broadcast newsroom workflows that let users access and share group-wide content remotely. www.bitcentral.com

Building4Media FORK Production

FORK Production features include ingest, low-res editing, Final Cut Pro integration, Media Asset

Management, new media, archive integration, newsroom integration and live assist playout. www.building4media.com

Cinegy Archive

An alternative to tape-based archives and digital archive and back-up solutions that require direct connections to video servers; relies on tiered storage model. Based on open standards and formats (AAF, MXF and MPEG2 and H.264), it lets loggers, story and video editors work on files in real-time at ingest. www.cinegy.com/jml/index.php/en/products-mainmenu-50/mam-menu/archive-menu.html

Crispin ArchiveManager

Manages the storage of digital programs on a mix of direct attached RAID drives, NAS storage and Blu-ray disc for long-term storage. www.crispincorp.com/ArchiveManager.html

Constellation VCM

See Line 1 Media, page 56

DAX Broadcast Archive

Media Asset Manager that supports Avid, Adobe, FCP, P2, XDCAM and RED and archives to LTO, Blu-ray, FTP/hard disk. www.daxarchiving.com/

Dalet Enterprise Edition

Media asset management for production and archiving with Web-based access to media archives including video, audio, graphics, documents and data files, such as EDLs and shotlists. www.dalet.com/MAM-for-Archiving

Dayang D3 MAM and iMAM

Media asset managers that control large broadcast collaborative workflows (D3 MAM) or smaller facilities (iMAM). iMAM features multi-format support, flexible catalogue module, and intelligent archive system and a user-friendly Web search service. www.dayang.com

DVS Spycer

Spycer lets users find, browse, edit and manage image data across a distributed content management network. Spycer detects all graphic file formats or file sequences, such as DPX, TIFF or TGA, as coherent clips. www.dvs.de

Digital Broadcast

Offers integrated software and hardware for automated media delivery transference, on-air verification, disaster recovery, archiving, and broadcast centralization.

www.digitalbcast.com

EditShare Flow and EditShare Ark

Flow and Ark extend EditShare's shared production storage workflows with advanced multi-channel ingest/browse on the front end and archiving/backup for a scalable, end-to-end media management workflow for post, DI and broadcast.

www.editshare.com

EMC Documentum Digital Media Archiving

Automates the management of content stored in federated repositories, lowering overall storage costs through the use of lower tiers of storage for historical assets and enabling intelligent repurposing and distribution.

www.emc.com

Empress Media Asset Management (eMAM)

Digital library system that holds and organizes videos, audios, images and documents.

www.empressdigital.com

Etere

Content management system for ingest, indexing, storage and retrieval of digital assets.

www.eter.com

FOR-A MediaConcierge Asset Management

Suite of tools for ingest, editing, accumulation, delivery, archiving and repurposing that support a diverse range of file formats.

www.for-a.com

Front Porch Digital DIVAdirector

Desktop- and Web-based media asset manager of real-time content, optimized for use with DIVArchive.

www.fpdigital.com

Front Porch Digital DIVArchive

Storage manager of disk and tape libraries.

www.fpdigital.com

Hamlet Reel-Check

PC-based software that features QC control for file-based media and live feeds and full library management and logging facilities for the broadcast and TV environment.

www.hamlet.co.uk

Hardata HDX Video

Stores, indexes, catalogs, searches, programs, and broadcasts audiovisual content.

www.hardata.com

IBM Content Manager

Manages all types of digitized content across multiple platforms, databases and applications by providing imaging, digital asset management, Web content management and content integration.

www-01.ibm.com/software/data/cm/cmgr/

Imagine Products ShotPut Pro and ShotPut RED

Offloading and management software for P2, Sony SxS, AVCHD, JVC ProHD and RED ONE that makes secure, auto-named copies of a card or hard disk's contents in up to three locations.

www.imagineproducts.com/index.php?main_page=product_info&cPath=1&products_id=2

Isilon SyncIQ

File-based replication for disaster recovery, disk-to-disk backup and distributed workflows.

www.isilon.com

Line 1 Media Constellation VCM Asset Management Software

Manages all content across tape-based or tapeless media workflows, from SAN storage, nearline storage to tape or optical archives with full search, retrieve and manipulation of video and project files.

www.line1media.com/constellation.html

Masstech Group INDIGO and TOPAZ

The company's revamped (as of March 2010) product line of Media Asset Management tools include these combination Media Asset Managers/Content Storage Managers that run on both Windows and Apple operating systems. TOPAZ is specifically aimed at smaller production and post facilities.

www.masstech.com

Media Proxy

Monitoring and logging system that records, archives, indexes and re-purposes broadcast and IPTV media and related information for broadcasters, media firms, and corporations.
www.mediaproxy.com

MediaBeacon R3volution DAM Suite

Secure Web-based media manager with drag-and-drop integration with Adobe CS3/CS4, Final Cut Pro and Microsoft office.
www.brightech.com/

North Plains TeleScope Studio

Web-based media manager aimed at small creative organizations such as ad agencies, broadcast studios and publishers.
www.northplains.com/

Open Text Digital Media Group

Artesia DAM and Open Text Archive Server DAM and archiving server for the broadcast, marketing and advertising markets.
digitalmedia.opentext.com/

SGL FlashNet and FlashBox

Content archive and storage management systems targeting the broadcast industry that use automated content-replication systems to provide synchronized, mirrored or like-for-like asset duplication, across same-site or geographically disparate locations.
www.sgluk.com

Squarebox Systems CatDV Production Management

Cross-platform logging, transcoding and database manager with integrated tape library.
www.squarebox.co.uk/

Sony Media Backbone

A new product category for Sony and introduced at NAB 2010, the open-platform Media Backbone product family grew out of the company's desire to address workflow issues created by the relentless uptick in 4K and 3D production and post. An open, frame-level asset manager tracks and manages digital files from ingest through to the deepest levels of the archive.
www.sonystyle.com

TMD MediaFlex

TMD modular software products manage production, delivery, archiving and re-purposing of media content in all formats including film and tape-based content to enriched digital media. MediaFlex manages both physical media and file-based content.
www.tmd.tv

Replication, Deduplication and Archive Software and Servers for SAN, NAS and Tiered Storage*

*This list does not include the wide range of storage options currently on the market. The following vendors offer some form of integrated archiving solutions or optimized architectures for third-party archiving software.

Active Storage XRAID

Disk-to-disk archival backup designed for post-production and broadcast and optimized for Apple Xsan and Mac OS X.
www.getactivestorage.com

Atempo Digital Archive

Features a drag-and-drop user interface that lets editors archive completed projects directly from primary storage—and retrieve projects archived by others. Comes with a plug-in to Final Cut Server that enables archiving to LTO tape directly from Final Cut Server.
www.atempo.com

BrightDrive

Storage servers for managing nearline archiving.
www.4bright.com

Bycast StorageGRID

Bycast StorageGRID software simplifies the management of massive fixed-content storage systems and lets organizations optimize their storage infrastructure and ensure the integrity and availability of their valuable data assets over their lifetime.
www.bycast.com

DataDirect Networks S2A9900

SATA storage with MAID infrastructure that uses 86% less power and 75% less rack space.
www.ddn.com

Digital Rapids StreamZ and StreamZHD

Real-time media encoding servers that works in tandem with video archiving appliances and storage servers.

www.digital-rapids.com

EMC Data Domain

Deduplication systems for disk backup, archiving, disaster recovery

www.datadomain.com

EMC DiskXtender

Migrates inactive data off higher-cost storage to lower-cost disk, tape or optical devices.

www.emc.com

JMR BlueStor

Enterprise-class servers for DI and color monitoring, direct-attached Fibre Channel PCIe RAID for online workflows and offline storage options for archiving and disaster recovery protection.

www.jmr.com

Omneon Media Application Server

Optimized to work with XenData's Archive series.

www.omneon.com

Overland Storage

Backup and recovery software and systems that target broadcasting, video surveillance industry and other corporate markets.

www.overlandstorage.com

Quantum StorNext

Data-sharing and archiving software integrates SAN file system and distributed LAN Client to centralize content and build cost-effective digital archives for long-term storage and content protection.

www.quantum.com

SGL/Copan Systems

Tiered-storage systems that work with SGL FlashNet.

www.sgi.com/copan

XenData Archive Series

Windows-based software using RAID cache and LTO tape storage.

www.xendata.com

Offsite Conversion and Archiving Services

Crawford Media Management Services

Atlanta-based Crawford Communications's new division (an offshoot of Crawford Post Production) specializes in high-bandwidth networks, post production workflows, legacy media formats, asset management, storage and disaster recovery.

www.crawfordmediaservices.com

Digital Migration Solutions

Services include digitization, encoding or transcoding of analog media, quality monitoring and consulting on long-term digital storage. Specializes in digitizing media for Avid and Final Cut workflows for Web sites such as iTunes and Hulu.

www.digitalmigrationsolutions.com

DuArt

Film restoration services include HD transfers of all legacy film and video formats and archival quality preservation.

www.duart.com

Haghefilm Conservation

Haghefilm Conservation is a division of Cineco Motion Picture Laboratories in Amsterdam, the Netherlands. It specializes in high-end conservation and restoration of film footage, using both conventional and state-of-the-art digital technology. Among its major clients: the George Eastman House in Rochester, NY.

www.haghefilm.nl

Iron Mountain

Offers cold storage of tape or disk archives in controlled environments.

www.ironmountain.com

Media Archive Systems

Offers encoding and content conversion systems, including warehousing of old and new media.

www.mediaarchivesystems.com/services.html

MediaKive

Founded in 2007 to digitize growing video tape libraries. Services include consulting on metadata migration, disaster recovery and asset management systems.

www.mediakive.com

Northeast Historic Film

Cold storage for film reels, videotape and still photographic material.
www.oldfilm.org/storage

Pacific Title Archives

Services include media asset management of digital files as well as inspection and climate-controlled preservation storage of analog media.
www.pacifictitlearchives.com/

ScreenSavers

Scene Savers is an archival services company that helps organizations preserve and provide access to their film and video collections. Services include film to video transfer, videotape remastering, digitizing, storage, digital asset management and consulting.
www.scenesavers.com

Technicolor Archive & Library System

Secure storage of physical and file based assets. Providing services for online, near-line and/or remote offline, complete with indexing, proxy and retrieval functions.
www.technicolor.com

VidiPax LLC

VidiPax provides extensive services including the traditional restoration of original video, film, data and audio elements; recovery of damaged or deteriorated originals; standards conversion; sonic and visual restoration; disaster recovery; system design and implementation; research and development; media collection assessment; archival services; cataloguing and data entry.
www.vidipax.com/

studiodaily

www.studiodaily.com