

The Cache-A Archive Appliance

Archiving and Interchange for a Post-Videotape World

If you're a professional in video production or digital film, you already know that managing change is key to staying on top of your game. Technical innovations, from acquisition formats to NLE features, have made possible creative capabilities that seemed far-fetched just a decade ago, and there's no reason to expect the pace to let up anytime soon. On the plus side, change opens new creative possibilities and brings increased productivity. But it also presents new challenges, requiring us to move on from solutions that worked well in the past but no longer fit our needs. In a file-based world, where video is data and videotape is fast giving way to cards and drives, one critical challenge is finding a new way to handle the storage and archival role traditionally played by videotape. The solution is the Cache-A archive appliance.

Developed by Cache-A Corporation, the Cache-A archive appliance combines the immediacy and random access of hard-disk storage with the longevity and reliability of IT-proven archival-quality LTO data tape, creating a high-capacity "data tape deck" that is self-contained, network-accessible, platform-independent, and extremely cost-effective. The result is a video-centric data storage appliance that makes an ideal archiving and interchange solution for post-videotape workflows:

- **Production** — Cache-A archive appliances streamline the creation of source masters from diverse non-tape media, protecting fresh footage while freeing cards for further shooting. And they allow fast multi-user access to captured video for QC, metadata editing, and other asset management tasks.
- **Post production** — Cache-A archive appliances not only archive finished projects but also solve the question of how to deliver those projects to the client. With Cache-A, masters, outtakes, source clips, and other associated assets can be stored together in a single 800GB cartridge (or spanned across a multi-cartridge set, if needed) from which any individual file is readily available.
- **Asset interchange** — Whether moving footage between facilities or between phases of production, archive appliances create compact, reliable video storehouses whose contents can be accessed and managed with drag-and-drop ease, pulled over any standard Ethernet LAN, and couriered around the world in a single-envelope.

While many data storage technologies on the market today can be adapted to handle some subset of the overall archival needs of professional video and digital film, the Cache-A archive appliance is the only

solution that addresses the full range of requirements, from reliability and retrievability for long-term storage to the practical workflow needs of today's fast-paced production environments. In this paper, we'll look at what those requirements are, how they are handled by archive appliances, and how archive appliances measure up compared to competing options. We think you'll agree that the Cache-A archive appliance provides a superior solution for the protection and exchange of valuable video assets.

The Post-Videotape Workflow

In recent years, two powerful technical advances have reshaped the way that video is acquired, prepared, and stored. One was the maturity of computer-hosted non-linear editors (NLEs) into devices that are capable not only of offline editing tasks, but also of full-resolution finishing, including effects and compositing. Instead of simply making it easier to edit from source tapes, the newer generation of NLEs meant that once the tapes from a shoot were digitized they were no longer required for completion of the project in post production. The workflow became file-based from ingestion onward, and the source masters themselves effectively became shelf-stored archives, to be revisited only if needed for later revisions or repurposing.

The second major development was the advent of tapeless professional acquisition with solid-state media such as Panasonic's P2, Sony's SxS, and RED CF or RED Drive. Tapeless camcorders mean that the video exists as a file from the moment of capture, which allows it to be loaded directly into NLE storage for post. But since NLE storage space must eventually be used for subsequent projects, and the expense of solid-state cards rules out retaining the cards themselves as source masters, the footage must also be transferred to an alternate medium so that a source master exists once the cards themselves are reused. What should that medium be?

The same question arises when we consider the need to pass materials between facilities that will handle different aspects of post production. And it's relevant as well after post, when the completed program, as well as all the associated source assets, must be delivered to the client and also archived for the future. To understand why the Cache-A archive appliance offers the best possible answer, we'll first identify the requirements for an ideal solution addressing these various situations, and then see how well archive appliances and other options each meet those criteria.

Defining Archival Requirements

As it turns out, there are quite a few bases that must be covered to meet the requirements of an effective system for video archiving, source master creation, and asset interchange:

- **Archival, not backup** — Backup and archival systems share the goal of protecting valuable data. But backup systems are primarily designed to protect the state of individual computers and/or their hosted drives in case of a failure leading to data loss. To do so, backup systems typically employ proprietary data formats, incremental backup, and other techniques that limit

universal access to the underlying data, generally requiring restoration of an entire volume to get at any individual file. In contrast, an effective archival system makes it as easy as possible for any networked user to access any archived file, both immediately and decades later.

- **Reliable** — Archival systems must keep valuable video assets safe, period. That means the system must be built around media that is not only robust enough to stand up to handling in the near term but is also specifically designed to provide long-term reliability as measured in decades rather than years.
- **File-workflow compatible** — With NLEs already allowing a file-based workflow from the point of ingestion, archiving tapelessly-acquired assets to videotape is a huge step backwards. An effective archival system must retain all the advantages of file-based systems, including random access and catalog/search capabilities that make it easy to identify and find source material.
- **Multi-user** — Today's workflows require that assets be simultaneously accessible to multiple users without time-consuming transfers. That means an effective archival system must connect directly to the facility's LAN rather than operating as local storage for a single host computer.
- **Self-contained** — Archival systems should be independent and platform-neutral, supporting file sharing protocols (e.g. NFS, SAMBA/CIFS, AFP, Bonjour, FTP, etc.) that address the entire universe of production computers. Systems built on computer-hosted drives are tied to a single machine type, thus depending on specific operating systems and drivers that may not be available when stored content is needed a decade or two later.
- **Interchangeable** — Effective asset interchange means that the media itself must be physically portable and its contents easily detectable and readable. Widely-adopted media with an on-board directory and data in a standards-based format allows interchange regardless of the application or software environment used during archiving.
- **Simple** — Archiving systems should be easy to integrate and easy to use. A system requiring specialized knowledge and/or additional personnel to operate and maintain is a drag on both productivity and the bottom line.
- **Scalable** — If your video business is thriving, the sheer volume of video you need to archive is constantly growing, especially if you work in high definition. An archival system should allow you to easily add capacity while keeping track of all content in an integrated system rather than in a fragmented patchwork that makes assets hard to find.
- **Future-proof** — Will hardware be available in the future to retrieve files archived today? The best protection against obsolescence is a platform that is widely-deployed, has a track-record of backwards compatibility, uses non-proprietary standards, and is specifically designed for long-horizon applications.
- **Cost-effective** — Cost shouldn't be an obstacle to archiving. Cost-effective systems allow you to err on the side of asset protection, understanding that picking and choosing what to protect will inevitably lead to regret down the road. Cost-effectiveness should be judged not only in terms of initial outlay, but should also consider factors such as media (\$/GB), operating efficiency, and personnel requirements.

- **DAM compatible** — Intelligent asset management can be an important aspect of an efficient video workflow, and archive systems need to be able to readily retrieve assets identified via Digital Asset Management (DAM) software. But because DAM systems themselves come and go, an archive system should also be sufficiently independent that DAM upgrades do not require complete re-cataloging of archived content.

The Cache-A Solution

Now that we have a clear sense of the attributes required for an ideal video archival system, let's take a closer look at Cache-A's self-contained Cache-A archive appliance to get a better sense of how it addresses these needs. The core of the appliance is its powerful, updateable system software, which intelligently integrates and manages a set of components that make up the hardware side of the overall solution. Those basic hardware elements are as follows:

- One or more large-capacity high-speed hard-disk drive (multiple drives may be RAID-configured).
- One or more A-Series (video optimized) LTO-4 data tape drive.
- Archival-quality, IT industry-standard LTO-4 data cartridges.
- Network interfacing for Gigabit Ethernet connectivity at greater than 50MBps.
- USB ports for P2 and SxS readers.
- Optional ExpressCard slots for direct transfer from P2 and SxS media.
- Optional SATA and SAS interfaces.
- Chassis housing the above elements.

By combining hard disk with data tape, the Cache-A archive appliance is able to offer the best of both worlds. The disk allows simultaneous multi-user access and provides a familiar paradigm for interacting with network-attached storage (NAS). A wide range of supported file transfer protocols, including NFS, SAMBA, AFP, Bonjour, and FTP, make the appliance compatible with (yet independent from) all commonly-used computing platforms. The inclusion of a hard disk also provides a staging area from which production personnel can handle pre-archiving tasks such as playing back footage for QC and adding metadata.

The archive appliance's archiving capability, meanwhile, is provided by the A-Series LTO-4 tape drive. The LTO line is the world's most widely adopted family of data tape drives, with over 2.5 million deployed to date for mission-critical applications in sectors including corporate IT, financial services, military, and government. Transfers between the hard drive and the LTO-4 drive occur transparently in the background at speeds up to 80MBps.

LTO-4 cartridges are designed to withstand the wear and tear of daily use and are certified for a 30-year archival life. Each tape holds at least 800GB, enough for more than 60 hours of 25Mbps standard definition (SD) or more than 15 hours of 100Mbps high definition (HD) content (actual capacity may be up to 1.6 TB using lossless compression). If a single tape isn't enough for a given project, the appliance will span the data across a multi-tape set.

File-based compatibility

Beyond extraordinary capacity, speed, and longevity, what distinguishes Cache-A data tapes from standard videotape is their compatibility with file-based workflows. A directory included on each tape includes file system information for every file (size, creation, modification, location, etc.). The directory can be viewed and read in Cache-A's own browser-hosted interface, allowing you to track and restore contents without additional archive management software. You can also quickly save files and folders to the appliance, as you would to any other generic data storage device, by using your computer's file system (e.g. Windows Explorer or Mac OS Finder) or by choosing File > Save from within an application.

The archive appliance not only provides access to the directory of the currently-inserted tape, but also maintains a catalog of all tapes that have ever been used on the system. Since each tape has both a user-defined name and a unique ID (which is also printed on the cartridge itself), the catalog lets you find the host tape of every file you have ever archived with the appliance. And if multiple appliances are used on the same network, their catalogs are seamlessly merged.

The catalog is viewed in Cache-A's browser-hosted interface, which allows access to files in three types of locations:

- **VTAPE** (virtual tape) — an appliance's own network-shared volume;
- **Client-shared volume** — a shared folder mounted by the appliance from any computer on the same network;
- **Direct-mounted volume** — a storage device that is physically connected to the appliance's ExpressCard slot or USB, SATA, and SAS interfaces.

Archiving to and restoring from files in any of the above locations is a simple drag and drop process within the archive appliance interface. Unlike typical backup systems, every file and folder is available independently; there is no need to restore an entire volume to access an individual file within it.

Archiving and interchange

Once files are archived with a Cache-A archive appliance, the destination LTO-4 tape becomes the source master, project master, or interchange media for the underlying assets, freeing acquisition media and NLE-connected storage for subsequent projects. For long-term archiving, the tapes can be efficiently stored for 30 years in a simple shelf library under standard climatic conditions. No expensive loader machinery or IT staff is required to quickly access the files should they be needed down the road.

With 800GB capacity in a compact form factor (approx. 4 x 4 x 1 inches), the LTO-4 offers tremendous data density, effectively eliminating physical space as a major constraint on library size. A single 6-foot by 6-foot wall rack with 14 shelves can store about 1000 cartridges, which works out to 60,000 hours of 25 Mbps SD and 15,000 hours of 100 Mbps HD.

Because each LTO-4 is stamped (both physically and electronically) with a unique ID, any archived Cache-A tape whose files have been accurately entered into a DAM can later be easily identified and located on the shelf without fear of confusion. Cache-A has also developed an API allowing integration with DAM and ingest packages such as Final Cut Server, CatDV, HDLog, and R3D Data Manager.

Cache-A tapes are similarly well-suited for project delivery and asset interchange. The LTO-4's robust cartridge and modest weight (about 10 oz.) make it ideal for couriating huge amounts of video to clients or to other production facilities. And because files are written to the tape in a "tar"-compatible format, anyone on the receiving end with an LTO-4 compatible data tape drive can retrieve all files by transferring to a hard drive and unpacking using a "tar" utility. Thus while a Cache-A appliance offers advantages such as the catalog and direct access to individual files, it is not required to recover the entire contents of a Cache-A archive.

Considering the alternatives

It's clear from the above that Cache-A's Video Archive Appliance meets all the necessary criteria to excel at source master creation, archiving, and interchange for video and digital film. And at current pricing of less than \$50 each, the media cost of LTO-4 works out to a mere \$0.06/GB, which makes Cache-A archives an unbeatable value as well.

Now let's consider how archive appliances stack up against the possible alternatives:

- **Videotape** — As noted above, the primary strike against archiving to videotape is that it reintroduces all the constraints of a non file-based medium into a workflow that has otherwise moved entirely to a file-based model. Among other issues, videotape content transfer is also slower than data tape and requires more QC to ensure problem-free transfers.
- **Blu-ray Disc** — BD capacity is limited to 50GB (DL), transfer rates are much slower than the Cache-A appliance, and the per-GB costs of BD-R and BD-RE are much higher than LTO-4.
- **Hard-disk only** — Removable Firewire and USB hard drives offer great short-term performance, and the price/GB has fallen considerably (though it's still high compared to LTO-4 cartridges). But hard drives alone can't meet the need for a solid archiving and interchange medium. They aren't robust enough to reliably endure the shock of being struck or dropped in transit, and they aren't designed to sit idly on the library shelf — they can freeze up if not used regularly, never to spin again. Even the best such drives are warranted for only five years, a fraction of LTO-4's 30-year lifespan.
- **Data tape only** — Standard LTO drives can be hooked directly to a computer and used with backup software to back up that computer's local data. But as noted earlier, backup and archiving are two distinct tasks, and systems optimized for one are less than optimal for the other. Backup systems are not primarily concerned with either long-term preservation or asset interchange, but rather with restoration of a computer if its hard drive fails. The backed-up data is typically tied to a proprietary software environment and a specific computing platform, and entire volumes must be restored to gain access to an individual file.

- **Tiered storage** — Large facilities can spend millions of dollars implementing complex, multi-tiered archiving strategies that progressively migrate assets to longer-term storage. The added costs of this approach include more equipment for additional layers of storage, more personnel for managing asset migration and maintaining equipment, and more operating expense for powering and cooling hard drives to store assets that could already be safely stored on archival-grade data tape. In contrast, Cache-A’s archive appliance offers both direct file access and immediate archival-grade protection, but does so with much greater cost-effectiveness and simplicity.

Cache-A Use Cases

By now it’s clear that when it comes to professional video and digital film, the Cache-A appliance is the only archival approach that fully meets the criteria for a complete archiving and interchange solution. And Cache-A’s utility isn’t limited to a single phase of the content creation process. The increasing popularity of solid-state acquisition means that the videotape-based workflow of “shoot/edit/archive” is being superseded by “shoot/archive/edit/archive.” Cache-A is uniquely suited to this new environment because it addresses not only the archiving challenges associated with both shooting (production) and editing (post), but also the associated tasks of asset interchange and asset management:

- **Cache-A for production** — In the field and in the studio, Cache-A archive appliances provide the most practical, safe, and portable solution for freeing high-cost premium storage devices for reuse:
 - Create archival-grade source masters of tapelessly-acquired footage, thereby protecting footage from the outset and avoiding later transfers from interim storage.
 - Directly connect most popular camera storage systems and cards to the appliance:
 - P2 cards and most P2 devices;
 - SxS cards and most EX devices;
 - CF cards and most RED devices.
 - Transferred footage can be immediately played back from the appliance, allowing QC before acquisition media is reused.
 - Self-contained, standalone appliance has a compact footprint for in-field operation.
- **Cache-A for post production** — In post, Cache-A archive appliances offer the most integrated networked solution for archiving while also providing a robust medium for delivery to clients:
 - Searchable catalog of every tape that has ever been in the appliance allows fast finding of files even on tapes that are not currently inserted
 - Editors in the workgroup can archive or access assets from any computer on the network.
 - Platform agnostic: barrier-free compatibility with heterogeneous workgroups including Mac OS, Windows, and Linux/Unix machines.
 - Workflow compatibility with Avid, Final Cut Pro, and other NLEs.
 - Cost-effective client delivery for dailies and finished projects (less than \$0.06/GB).

- **Cache-A for asset interchange** — LTO-4 tapes created by Cache-A archive appliances offer a reliable, self-describing medium for the interchange of assets between facilities at all stages of the content creation process:
 - 800GB capacity allows a single LTO-4 to move vast amounts of video, audio, and graphics as well as other associated project files.
 - Robust cartridge protects assets during handling and transit.
 - Light weight and compact size makes shipping easy, cost-effective, and competitive with transferring huge amounts of data over dedicated data lines.
 - “tar” compatible format makes tape contents accessible to any facility with an LTO-4 drive.
- **Cache-A for asset management** — Cache-A archive appliances feature integrated cataloging capabilities that aid asset management, and they also work hand-in-hand with dedicated asset management systems to provide a complete solution that includes physical transfer to LTO data tapes:
 - Cache-A’s searchable catalog maintains user-added metadata (e.g. keywords) as well as associating unique barcode and location information with each tape.
 - Information for a given tape’s files can be searched even when the tape itself is not mounted.
 - “tar” compatible format future-proofs the system and makes archive contents retrievable without a Cache-A system.
 - Cache-A API facilitates deeper integration of archive appliances into DAM and ingest workflows based on packages such as Final Cut Server, HDLog, CatDV, and R3D Data Manager.
 - Assets can be easily found in a Cache-A catalog using a DAM-provided file path.
 - Cache-A’s independent catalog allows changing your DAM without obsoleting your information about archived assets.

Conclusion

As we’ve worked through the challenges of archiving and interfacing for today’s video and digital film industries, we’ve seen that the Cache-A archive appliance is the only solution that fully addresses the needs of a post-videotape world across the entire production workflow. No other system combines file-based workflow compatibility, multi-user network access, searchable cataloging, and robust physical interchange. No other system offers comparable simplicity, control, and cost-effectiveness. And no other system protects your irreplaceable assets with the reliability and longevity of 30-year archival-grade media while allowing right-now access to any individual archived file. For no-compromise video and digital film archiving, there is only one choice: the Cache-A archive appliance.