

LTFS Continues to Hit the Mark in Media & Entertainment



LTFS TODAY... A DECADE OF TRANSFORMING DIGITAL MEDIA WORKFLOWS.

Written by



LTFS Continues to Hit the Mark in Media & Entertainment

Introduction

This study of LTFS in the media industry was originally published in 2012. Much of the information presented at that time covered the potential uses of what was then a radically new file system for managing data on tape. Considering it has been nearly a decade since LTFS was introduced, it is now an appropriate time to review its status in the industry.

In 2012 we noted, “In the media and entertainment industry, content is king. The majority of this content is now produced in digital form and virtually all of this content has digital distribution; that content is now digital data. Protecting that content, the lifeblood of this industry, with the right data storage solution is more important than ever.” That statement is all the more true today. The amount of content is significantly greater, the size of content objects is larger, and the vulnerability of content to theft and cyberattack has only increased. Safe, secure, and efficient digital storage of media content remains a critical industry requirement and practice.

Digital content continues to grow rapidly due to the demands for more data-intensive formats. Then, it was HD, 2K, with 3D and 4K coming. Now, it’s 4K, HDR, 3D, IMAX, and complex audio, with 8K coming and ever more digital effects. Storage for this digital content can be expensive and cumbersome, adding to soaring production costs. Flash, disk, tape, and cloud are the most prevalent storage technologies, each with its own set of trade-offs.

Since the introduction of Linear Tape File System (LTFS) in 2010, the media and entertainment industry has taken full advantage of LTFS being an open-source, self-describing format interchangeable standard with ease of use via most operating systems (macOS, Windows, Linux). One major advantage LTFS has provided for all industries is the interchangeability of content on a robust portable format such as LTO tape. LTFS and associated products have already been recognized by the industry, winning several awards, culminating with a prestigious Technical Emmy Award. LTFS has become a SNIA Specification and an ISO Standard.

This edition of the whitepaper revisits and updates how LTFS provides significant advantages for media and entertainment applications. It includes an overview of LTFS updates, evaluates the pros and cons versus other storage technologies, and takes a look at how LTFS and LTO tape in general enhances media industry workflows and processes.

Table of Contents

- Notes to the Reader2
- What is LTFS?2
- Why was LTFS developed?2
- How does LTFS work?3
- What are the Benefits of LTFS?4
- Where does LTFS fit into media workflows?.....6
 - LTFS in Production8
 - LTFS in Post Production.....8
 - LTFS for Distribution.....9
 - LTFS for Digital Media Archiving9
 - LTO with LTFS or Cloud Storage for Digital Media Archiving?10
- Is LTFS the answer by itself?12
- LTFS Considerations14
- LTFS Implementation Options15
 - General LTFS Applications.....15
 - LTFS Products for Media and Entertainment15
- What Might the Future Hold for LTFS?17
- References17
- Appendix19
 - Storage Tiers19
 - LTFS Success Stories:19

Notes to the Reader

- Both LTFS and LTO-5 were introduced at the same time in 2010. Since 2010, several more generations of LTO tape have come to the market (LTO-6, 7, and 8), with more to come. LTFS has and will work with all these follow-on generations of LTO. Consequently, for the purposes of this review, “LTO” will be a reference to Generation 5 onwards.
- Digital media content is often compressed either in capture/production, post production, or distribution processes. Typically, compressing previously compressed data provides no benefit. For this reason, the examples in this review will reference the uncompressed capacity of LTO-8 tape formatted for LTFS as 11.4 TBs (95% of 12 TBs) and LTO-8 tape drive bandwidth as 360 MB/second.

What is LTFS?

Linear Tape File System (LTFS) is a multi-vendor open format specification for storing data files and an index of those files together on data tape. It relies on the support of the tape technology for partitioning that was introduced in the LTO technology roadmap with the LTO-5 format specification and carried forward in subsequent generations. The metadata stored in one partition provides an index or directory of the data stored in the other partition.

The Storage Networking Industry Association (SNIA) (snia.org) manages the LTFS Format Specification. The current version of the format, Version 2.5, was released in May, 2019 (https://www.snia.org/tech_activities/standards/curr_standards/ltfs). The LTFS Format Specification has also been adopted as an ISO/IEC Standard - ISO/IEC 2019:2016 (<https://www.iso.org/standard/69458.html>)

An LTFS driver is software that extends operating system functionality, enabling the data on an LTFS tape to be accessed by a file system similar to that found on disk drives using commands like copy, move, and delete. Once LTFS is installed, applications such as file explorers or finders can access data on LTFS-written tapes much like disks and USB devices. This includes operating system supported user friendly drag and drop interfaces.

The combination of the tape format and operating system drivers is LTFS.

Why was LTFS developed?

LTFS was developed to make “...using tape as easy, flexible, portable, and intuitive as using other removable and sharable media, such as a USB drive.” Tape for data storage is an established, trusted, and proven technology. Although the tape “model” hasn’t changed dramatically over the years, the speed, storage density and features of data tape has improved significantly. Uncompressed maximum data transfer rates were 120 MB/s for LTO-5 and are now up to 360 MB/s for LTO-8. More dramatically raw storage capacity has risen 800% from 1.5 TBs in an LTO-5 generation tape to 12 TBs uncompressed for LTO-8. The LTO Roadmap show the promise of continuing the capacity growth trend into the future. LTO tape remains the most reliable and inexpensive commercially available storage medium.

To enable easier use of this reliable and efficient storage media LTFS was designed to present a standard file system view of the data stored on the tape. With LTFS, accessing files stored on the LTFS

formatted media is very similar to accessing files stored on other forms of storage media, such as disk or removable USB flash drives. LTFS for data tape is analogous to how time-code systems are used in video and sound editing to search for individual clips or frames.

While LTFS is useful for a wide range of applications, it particularly appeals to those working with unstructured data and digital media. This is especially true for the media and entertainment industry with its numerous and large files (i.e. video, digital cinema, audio) to manage and preserve. The high data transfer rate and dense cost-effective storage of LTO tape are appealing to the industry, but were somewhat difficult to use... until LTFS.

How does LTFS work?

LTFS makes data on tapes accessible in a manner like accessing data on disk drives or USB flash drives. These devices contain blocks of data as well as an index (e.g. FAT32-File Allocation Table), which holds metadata information about a file (e.g. name, size, creation date, extended attributes, etc.) and a map of which blocks of data belong to which file.

Prior to LTFS, tapes contained only blocks of data and file marks in sequential order, without a directory to locate a specific file. External applications such as backup and archive storage management software maintained the mapping of which blocks belonged to which files. If that mapping were lost, the data on the tape would be present but not easily usable.

Due to the serial nature of tape, there is no inherent “in place” update capability. If a block is rewritten, all subsequent blocks must also be rewritten.

Alternatively, updating an index at the current end of a tape is not practical as it would move to a new physical location every time data is added.

But with LTFS, tape doesn’t have this limitation. Beginning with LTO-5 a standard LTO tape cartridge can be formatted for LTFS. This segments the tape into two partitions: one for the data and one for the index, so the index partition can be modified as needed **without** affecting the data partition.

LTFS at a Glance

An LTO tape cartridge is divided into 2 partitions

- One partition contains an index
- One (larger) partition holds the data

The index partition can be updated without affecting the data partition

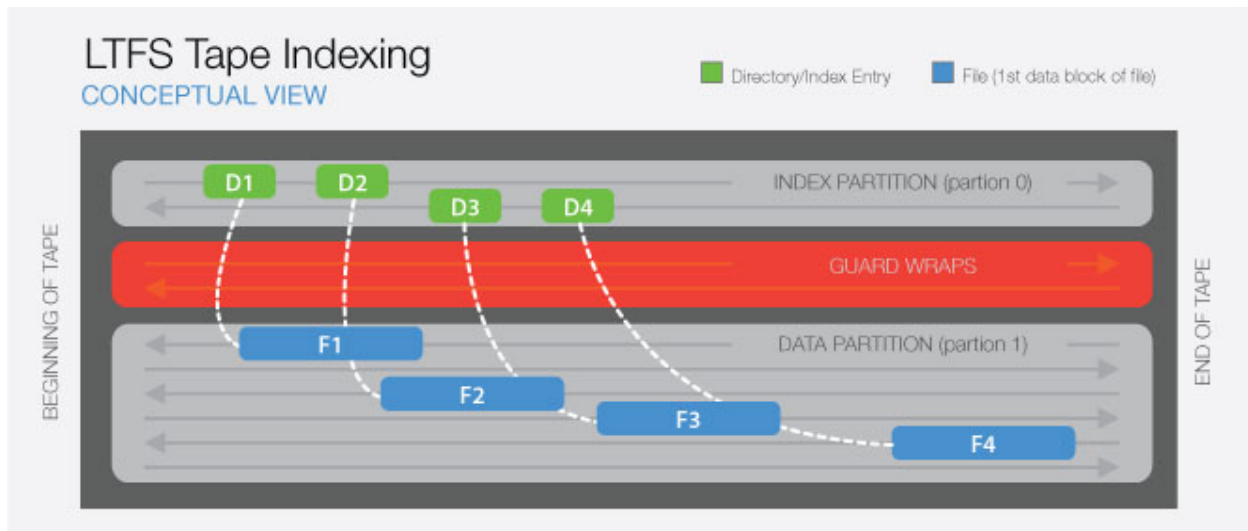
- Analogous to disk drives, USB flash drives, etc.
- Index data is cached for faster operation

Human readable XML index file maps which blocks belong to which files

- Parsable by software or manually
- Minimizes dependencies and preserves open access

Operating system drivers

- Data on tape appears as a standard file system
- Data on tape is usable by generic applications which may be unaware they are accessing data on tape
- Provides the ease of use similar to other storage devices



The combination of the partitioned LTO cartridge, format of the LTFS index, and conventions for the index placement is the magic behind LTFS. This creates a self-describing tape where a user can see the tape cartridge and its contents in the operating system directory tree browser (e.g., Windows Explorer, macOS Finder) and can copy, paste, or drag and drop files/folders to and from the tape. LTFS also supports extended attributes which enable custom file metadata.

Within the partitions, the tape contents are still stored via the trusted decades-old format of blocks of data and file marks. It is the interpretation of the contents of the blocks that is LTFS, **NOT** some radical alteration of established tape technology. LTFS is as remarkable for what it is not as for what it is.

The LTFS specification defines the format and layout of the tape, but not how it got there. Tape system vendors supply LTFS drivers for various operating systems which read and write according to the LTFS specification. Anyone is free to create their own software to utilize the LTFS specification instead and some have. Several tape system vendors provide LTFS verification utilities and the LTO Program has an LTFS Compliance Verification process. Further information is available at www.lto.org/technology/ltfs/ltfs-compliance-verification

What are the Benefits of LTFS?

LTFS combines the traditional advantages of tape (reliability, longevity, portability and low cost) with a familiar 'file/folder' type interface providing users faster easier direct access to files on data tape.

OPEN SPECIFICATION

The LTFS format specification and implementation scheme are in the public domain. It was adopted by the LTO Program group. It is now a SNIA Specification as well as an ISO standard. The LTFS Specification is independent of the tape media as long as the media supports partitioning for Reads and Writes. All LTO tape system and LTO tape media vendors support LTFS. Keeping user data stored in an open format has the advantages of allowing for multiple providers, increased options for users and interchange between both competitive and complementary offerings.

INTERCHANGE AND PORTABLE

LTFS tapes contain an index which describes the contents. There is no required connection with the software system or application that wrote the tape. This differs markedly from proprietary backup and archive storage management applications which usually require special tape formats

or metadata transfers from the application to provide access to data on the tape. In this regard, LTFS tapes and portable drives (HDDs or SSDs) are equivalent in portability.

EASY TO USE

LTFS is an extension of the operating system, so now tape is easy to access in a manner like accessing disk drives, USB drives, etc. A user only capable of using a mouse to drag and drop files can now be a tape user, heralding a new era of accessibility.

Numerous applications will work well without changes. Note that while an LTFS tape *appears* to the system just like a disk, it still *acts* like traditional data tape with its latency and serial access.

MULTI-PLATFORM SUPPORT

To complement the open specification of LTFS and the portable self-describing characteristics of an LTO tape written using LTFS, most LTO tape system vendors provide LTFS versions for all three dominant operating systems used by the media industry: Linux, Windows and macOS. This further enhances the benefits of interoperability, portability, and ease-of-use.

RELIABLE AND ROBUST

The LTFS specification contains a number of features designed specifically to enhance data recoverability under various circumstances. Multiple copies of the index are stored in both the index and data partitions. Older copies of the index are retained so the tape can be “rolled back” to a previous state with use of LTFS functions. The index format is in a human readable XML format and well documented in the specification.

SECURE

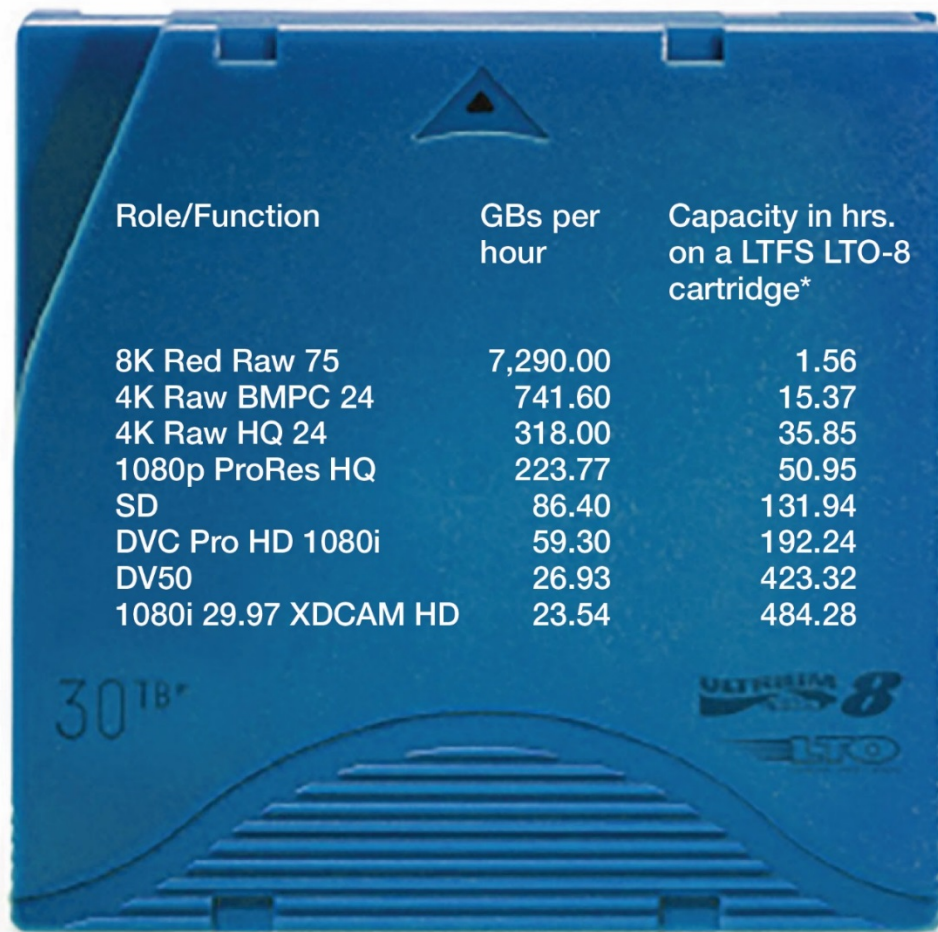
LTFS adds to LTO tape reliability features, which include tape drive hardware-based encryption to secure sensitive information and read-after-write verification to help ensure data integrity. Once an LTO LTFS formatted tape cartridge is removed from a tape drive/library it becomes a fully ‘air-gapped’ self-describing, reliable data repository. Current versions of the LTFS Specification also support implementation file permissions.

ECONOMICAL

LTO tape remains the most economically and the most reliable data storage media. This is why a significant portion of cloud storage providers utilize tape in their storage of long-term data. LTFS can further lower the cost of accessing and managing data on tape.

The economic benefits of LTO with LTFS can be particularly evident for the prosumer, small organization or budget conscious professional user with large long-term data storage needs. Stand-alone LTO-7 tape drives are currently available from \$3,000.00 to \$4,500.00 depending on method of attaching to a Windows, macOS or Linux workstation (i.e. SAS or USB 3.0). The LTFS stand-alone tape drive software drivers are free and many vendors provide basic software interfaces for utilizing the tape drive. With that investment all the other LTFS and LTO benefits are available.

A fact of life in media & entertainment is the demand of ever-increasing visual resolution, audio complexity, quality, and formats (HDR, 8K, Dolby Atmos, etc.). This creates more files and larger files, all of which must be managed. Keeping hours of such media online on disk quickly becomes cost prohibitive. This is especially true with uncompressed production-quality content as the following chart dramatically illustrates. LTO tape, especially with LTFS, can address the challenge.



Role/Function	GBs per hour	Capacity in hrs. on a LTFS LTO-8 cartridge*
8K Red Raw 75	7,290.00	1.56
4K Raw BMPC 24	741.60	15.37
4K Raw HQ 24	318.00	35.85
1080p ProRes HQ	223.77	50.95
SD	86.40	131.94
DVC Pro HD 1080i	59.30	192.24
DV50	26.93	423.32
1080i 29.97 XDCAM HD	23.54	484.28

30TB⁺ ULTRIM 8 LTO

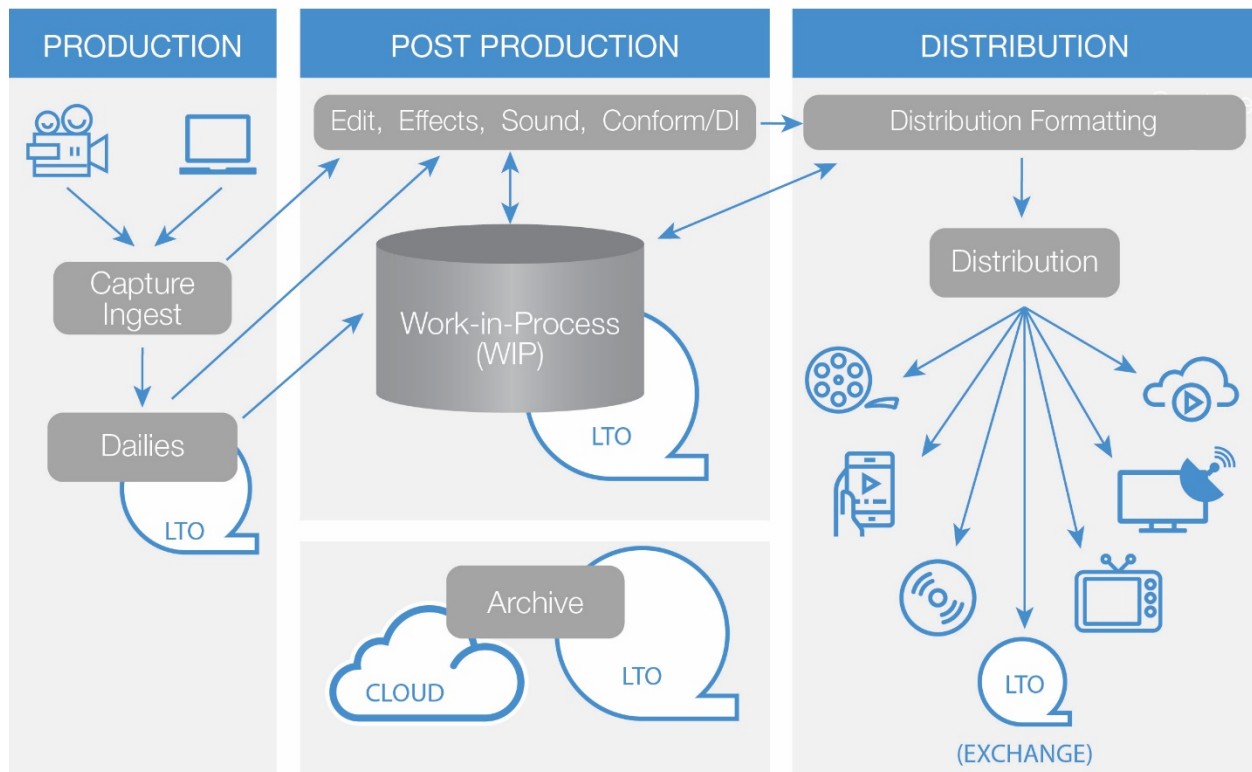
* Assumes 95% efficiency, 11,400 GB capacity

Where does LTFS fit into media workflows?

LTFS opened up new options for digital content storage and distribution in media and entertainment workflows, providing lower-cost higher reliability alternatives.

The diagram below presents a high-level view of a media workflow from production/capture through distribution. Before LTFS, LTO tape had been used primarily for archive and backup, including the backup of content created at Capture/Ingest. Now with LTFS, LTO tape is a primary interchange medium and provides a less expensive and easy-to-integrate working storage option for any of the functions not requiring low latency access to the content.

The Media & Entertainment File-based Workflow



Role of LTO/LTFS in a Media & Entertainment File-based Workflow

Role/Function	Long-Term Reliable Access	Lowest Total Cost	Open Format Interchange
Production	✓	✓	✓
Post	✓	✓	✓
Distribution	✓	✓	✓
Archive	✓	✓	✓

With the advent of digital technologies for capturing moving image (video and film) and sound, content producers and distributors have transitioned from analog/linear workflows based on film or videotape technology to digital/non-linear workflows based on the manipulation of data files. Historical analog workflows relied on physical video tapes and film elements as the medium for moving content from one process to the next. In the digital workflow the file is the unit of work being moved, whether by network file transfer, removable disk drive, or LTO tape.

The benefits of LTFS make it an excellent fit in a number of common file-based media workflow use cases. Allowing an LTO tape cartridge to be easily accessible with LTFS through a file/folder interface integrates into many file-based workflows.

LTFS in Production

LTO tape is firmly established in the media production environment with the advent of digital media. Historically, this has been the result of a mandate to many feature motion picture productions by insurance companies that content (data) captured on set or on location be archived to LTO tape on a daily basis. LTFS helps facilitate that process by being easier-to-use and more robust. The open and self-contained LTFS format is useful if the tapes are to be sent offsite, archived, or shared with a variety of recipients. A number of vendors provide on-set solutions and many of these use open LTFS format tapes. For production LTFS efficiently supports several important requirements:

CAMERA MEDIA OFFLOAD AND REUSE

Digital cameras encode motion images directly to solid state devices (SD, P2, etc.) or removable disks (e.g. XDCAM) in the camera. These media cards are quite expensive. Fast transfer of their contents to LTO tape with LTFS enables reuse of this expensive media, reducing the number cards or disks that must be purchased or rented.

BACKUP

Backup of daily footage to LTO tape is a common requirement as the loss of a day's worth of production is very costly. LTFS facilitates backup by enabling small portable independent systems to easily write daily content to tape.

TRANSPORT

The density and cost of LTO tapes with the self-describing capabilities of LTFS combine to create a very effective transport medium. Large amounts of data might be sent more quickly and economically than network-based transmission methods. This is especially compelling for digital productions which can produce terabytes of data for every day of shooting.

ARCHIVE

LTFS formatted tapes can be easily imported into an LTFS compatible archive by simply reading the index and adding the file metadata to an archive management system's catalog. The transport media and the archive storage media are one and the same under this scenario. The "import bandwidth" of tapes being added directly to a library en masse *far* exceeds any solution which requires movement of the actual data.

LTFS in Post Production

Post production is where all the content elements of a project come together: camera-captured content, digital effects, graphics and sound. It is in the post production process where all these elements are repeatedly edited, enhanced, assembled, and reviewed. A production is a collection of data files and post production has become a very technologically challenging part of the overall media workflow. Post Production uses medium to very high-resolution content (very large files), multiple processes and often multiple companies and departments working on the project and therefore the files. Making the content readily available is both an economic, logistic, and security challenge.

In an ideal post production storage scenario all the content (files) for a production (film, TV show, commercial, etc.) would be stored in a high-performance shared storage system with access control and integrated backup. The storage environment would ideally be accessible by all the post production functions and vendors. The reality is many of the functions and vendors will not have direct access to such a shared storage system. Files will have to be moved either by network or by physical media

(removable drives or tape). LTFS, available for the key platforms in media production – Linux, Mac and Windows – provides an excellent scalable and economic interchange method.

Another reality in post production is the ongoing need to manage the shared storage systems. The data files for a film project can easily need 10's if not 100's of terabytes of storage during post production. Multiple projects may be active at any one time and it is not unusual for project priorities to change – delays, postponements, and rush jobs. Again, economics typically drive a tiered storage solution utilizing a combination of high performance disk or Flash for online data, moderate performing lower cost disk storage for nearline and tape for work-in-process active archive storage, portability, and interchange.

While tape has always been a natural choice for the offline tier, LTFS enables tape to be a good contender for the nearline. LTFS enables applications developed for the traditional nearline disk tier to access data on LTFS tape without being aware the data is actually on tape. Latency (but not bandwidth) is the key performance metric which determines suitability for a particular use case.

In one popular scenario files are stored inexpensively on LTFS tape and brought online “just in time” for sharing, editing and further enhancement on disk. Once this work is completed, the new or modified files are written back to tape, leaving expensive online storage available for newer active files. Each technology is efficiently utilized where its cost and benefits are most effective.

Bandwidth “By the Box”

Have to move nearly 300 TBs of data fast? LTO tape formatted for LTFS functions remarkably well as a data transmission and exchange technology.

The US Postal Service will deliver their large flat rate box holding 26 LTO-8 tapes in their cases (about 296 TBs of capacity using LTFS) between major US cities for \$19.95 in 2 days. To move that much data over a network in 48 hours would require a 13.7 Gb/sec link running at 100% efficiency.

Need that data there tomorrow morning? Expedite it: the box can go from Los Angeles to New York in under 15 hours for less than \$220. You would need an end-to-end network faster than 45.4 Gb/sec to beat that box. (ship LA to Paris in 2 days for \$382.) With sophisticated technology like larger boxes, “bandwidth” can be scaled even higher.

To secure data in transit, LTO provides standardized hardware-based data encryption. Software-based encryption is also an option.

LTFS for Distribution

While network-based solutions are used for most business-to-business (B2B) distribution needs, they do require significant network bandwidth, technology, and expense.

LTFS formatted LTO tape provides the industry a common dependable exchange medium for large digital assets and large collections of digital assets. Its open specification and cross platform interchangeability distinguish it from other options. Add to that the storage density, reliability and robust nature of LTO tape and LTFS is an option for B2B media distribution.

If security during distribution is critical the content on an LTO LTFS tape can be encrypted. Encryption can occur before files are written to tape or via standard LTO tape drive hardware-based encryption which is compatible with LTFS.

LTFS for Digital Media Archiving

The need for highly predictable and reliable storage technology to support digital archiving is even more important today as growing amounts of valuable content is being captured, produced and distributed in digital formats.

The archives of a media company contain many of the company's most valuable assets. MGM and Miramax film library sales being historic examples and Disney obtaining 20th Century Fox's film archive as a significant asset in the Disney Fox merger being a current example. The archive masters of films, TV shows and music have real economic value. A recent lawsuit against Universal Music Group regarding the destruction of master recording during the 2008 Universal Studios lot fire alleges losses of up to \$150 Million.

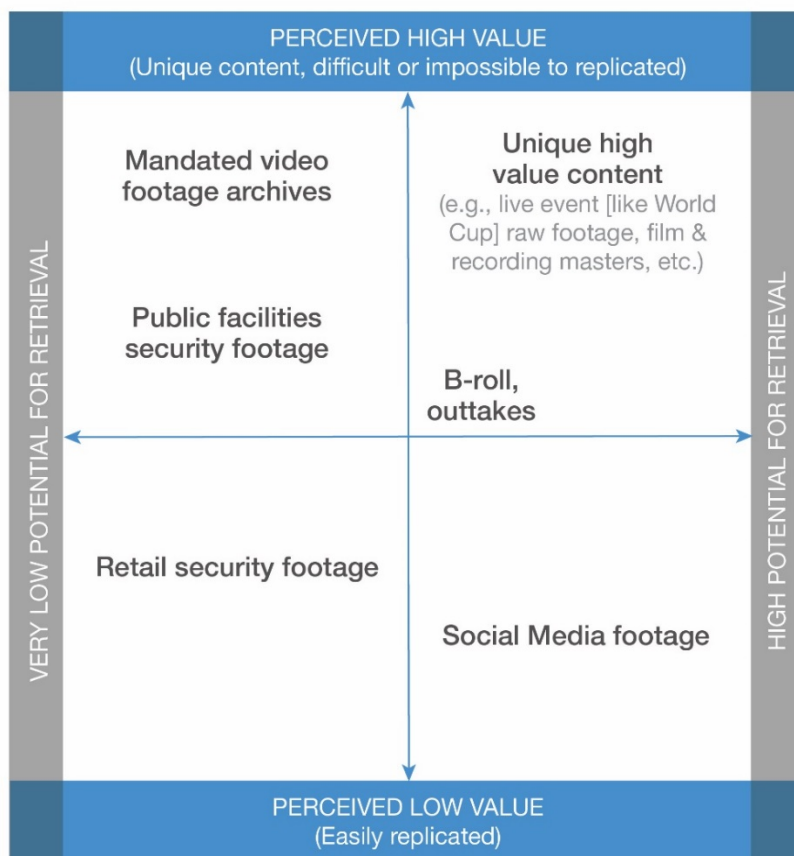
Historically, studio archives were analog film and recorded magnetic tape. Now, the archive of a project will likely only exist as digital files.

Data tape has long been the technology of choice for digital archiving. The overriding quality metric for a successful archive is whether the data can be retrieved when needed. With the addition of LTFS, tape can now provide ease of locating and retrieving files. LTFS was designed to have superior recoverability versus other IT storage technologies. LTFS minimizes dependencies, maximizes recoverability and facilitates the use of LTO tape for long term archives of high value data.

LTO with LTFS or Cloud Storage for Digital Media Archiving?

The answer is ... it depends. It may be both storage methods, or one or the other, or in some rare circumstances neither. There are many factors in determining the best digital media archiving strategy for your organization's assets. One approach is to evaluate an asset or asset class on two key characteristics: the perceived value of the asset and then how often will that asset be retrieved.

Content Archiving – Value/Retrieval Quadrant Diagram

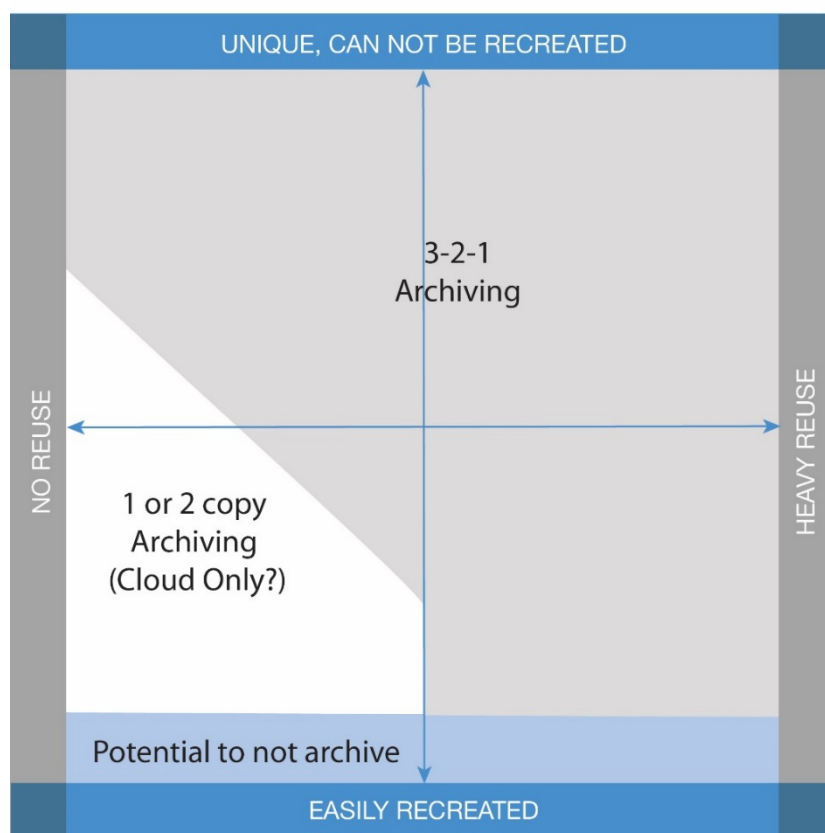


The following diagram uses these two characteristics in a quadrant chart providing one approach for evaluating archiving needs using some representative digital content examples.

The value of a digital media asset is as much a qualitative assessment as a financial one. The initial capture of a live event – raw camera files of goals in a World Cup final, the master digital elements of the latest superhero movie, the digital master files of a pain-staking restoration of a film classic or recording session, or live performance – is the starting point for the value stream of an asset. Recreating events like these to capture them again would be extremely expensive if not simply impossible. Of course bit-for-bit digital copies of these original

captures can and should be made and archived using best practices. The original capture, highest-resolution versions of the asset and its elements need to be protected, archived, and tightly managed to preserve their value. The versions made for distribution and revenue generation will almost always be derivatives, compressed to meet the requirements of the distribution and delivery networks.

Digital Media Assets Archiving Quadrant Diagram



Conversely, if a digital asset can be easily recreated or acquired at minimal cost does it really need to be archived. Even where there is limited perceived value, there may be legal or regulatory requirements to maintain certain classes of digital media for periods of time.

The other key variable is the potential for retrieval. Many productions create new show content from past shows or past events. Examples would be reality, talk, and news magazine shows. Fast predictable access to past content is essential. It can be difficult in most productions to forecast when and how much content might be retrieved.

Keeping in mind the digital asset groups characterized on the quadrant chart the various digital archiving approaches can be overlaid on the quadrant.

Assessing your classes of digital media assets on these two key variables helps point to appropriate digital archiving strategies. In the majority of cases following the classic '3-2-1' digital archive strategy is the best practice:

- Make at least **3** copies of the asset
- Store on at least **2** different types of media (we would recommend at least one copy be on LTO with LTFS).
- Ship at least **1** copy (probably the LTO with LTFS copy) to a location geographically distant enough to avoid any catastrophic event that might befall the primary location. Store that copy offline (on the shelf) to avoid hacking or ransomware risks.

In the '3-2-1' approach one copy could be on cloud storage. However, given the pricing models of most cloud storage providers, try not to retrieve content from that copy. Most cloud storage pricing models make it quite attractive to store content. However, retrieving some or all of that content for review or reuse can be quite expensive. Even a modest percentage of retrieved content can generate charges that outweigh the base monthly storage charges. Worse yet, the retrieving costs can be difficult to predict.

For some classes of digital assets with low perceived value and very infrequent access (the lower left-hand area of the diagram) a dual or single copy archiving strategy may be appropriate. A 'cloud storage only' approach maybe sufficient for those 'just in case' scenarios or where there is a mandated archive requirement.

Finally, there is some amount of captured content that doesn't need to be archived. The question many organizations confront is, is it cheaper to spend the time reviewing content to determine if it needs to be archived or to just save it all? This is an area where automated intelligent content analysis might help.

It is interesting to note that if your organization uses a cloud storage service for archiving, the odds are that at least one copy of that asset stored on a cloud service will reside on tape. As a group, the cloud storage service providers are very large users of LTO tape.

Is LTFS the answer by itself?

It can be for small or straightforward use cases. For many situations however, driven by the archive size and use cases, additional software will be needed. Different users will want different levels of management for their solution, from automated and policy-based managed HSM (tiered storage) to a stand-alone LTO drive with LTFS and simple drag and drop.

In a single drive LTFS implementation, the tape already contains an index viewable through Linux, Windows, or Mac utilities. In addition, the user has drag and drop access. However, that index only covers the contents of the tape cartridge currently mounted in the tape drive. There is no centralized or library-level index that contains the content on multiple tapes. For smaller scale implementations it is practical for a user to individually track the tape contents (e.g. maintaining a separate spreadsheet or data base) and receive the benefits of LTFS.

For larger storage environments, implementations of LTFS are available for tape libraries. These applications map the contents of a tape library or library partition as a set of directory trees, typically one for each LTFS formatted tape.

The LTFS tape format is unchanged and identical to the single tape implementation. Index data are typically cached for rapid access and tapes are loaded as needed to access data. These implementations may present as a NAS or an object store.

LTFS Applications for other industry verticals

LTFS is well suited for storing large unstructured data and large collections of unstructured data. The Media industry is just one of many with very large unstructured data storage requirements. Other examples include:

Healthcare – Medical applications such as imaging, X-ray, and ultra-sound now feature increased resolution and motion imaging. Regulations and best-practices call for secure long-term storage solutions for diagnostic, research, and legal purposes.

Security – Video surveillance is generally not high bandwidth but the sheer number of parallel channels (it can be thousands for a single facility such as a casino or airport) and the duration of each feed produces a tremendous amount of data. LTFS facilitates accessing the right file. Furthermore, a workstation can be used to review footage independent of the system that recorded it.

Biological Data – Human genome mapping produces enormous amount of data - 3 billion base pairs per human! Researchers leverage these massive data sets to examine correlations between genetic mappings and specific diseases. The portability and accessibility of LTFS are well suited as is LTFS's seamless support of LTO technology's lossless data compression.

Cloud Storage – Cloud storage accumulates vast amounts of data. LTFS formatted tapes provide a very economical and easy-to-use method for transferring large amounts of data (even between disparate operating systems). The same tapes used for a transfer can be retained to provide a local offline restore capability.

Many media asset management systems (MAMs, DAMs, PAMs, etc.) and media content-oriented archive and storage management software solutions now support LTFS; if not as the native tape storage format, then as an import/export supported format.

Media companies gain content accessibility and operational flexibility with LTFS-enabled backup and archive storage management solutions. The storage management application is **no longer required for unlocking data in the archive**. Any one of the LTFS formatted tapes can be pulled and its contents read at a workstation with a stand-alone LTO tape drive using LTFS. Additionally, should there be a catastrophic failure of an LTFS-based storage management application database, the contents can be rebuilt by reading the index partition of each of the managed tapes. Similarly, data stored on LTFS tapes can be moved between storage management products from different vendors providing more options than proprietary formats.

Tape remains the medium of choice, or even the only choice, for many very large storage applications. While cost per gigabyte for tape and disk continue to shrink, tape maintains the advantage over disk when the overall total cost of ownership is taken into account. LTO tape provides the high bandwidth and economical storage density required, and LTFS makes it easier to locate, access and share the files stored on tape.

The point at which tape is more cost effective than disk is not based just on hardware cost. Complexity and the associated cost of managing tape also figures prominently in the decision-making process. LTFS reduces tape technology's "barrier to entry" and lowers the cost of tape-based solutions. Tape becomes more attractive when you also consider factors such as disk mirroring and snapshots and of course the total-cost-of-computing (e.g. power and cooling).

What to look for in an Application Supporting LTFS

When evaluating an application supporting LTFS, consider the following:

How do users access the files on the LTFS tapes?

- Are they shared out via a NAS-like or object store appliance or is custom interface required?
- How easily can your existing software and workflow integrate?

How well does application/appliance scale?

- What is the prudent upper limit on the number of files it can manage and still meet performance metrics?
- Does it mediate and manage multiple concurrent user requests?

How do users locate their files?

- Do they search on metadata fields or just by filename?
- Do they select individual files? Contiguous groups? Arbitrary sets?
- Do they specify sets of files via patterns or lists?

How efficient are the data transfers?

- If there are dozens or thousands of files being recalled, can it maximize tape streaming? Does it use the location of the files on tape to sequence reads?
- What is the net bandwidth for various recall scenarios?
- Can it write data to optimize read patterns?

Does the application support the current version of LTFS and new capabilities?

- 2.5 is the current SNIA specification

LTFS Considerations

Many of the characteristics of LTFS are associated with the physical tape it resides on and are primarily related to latency. To access data on a tape, the tape must first be loaded into a drive. If all drives are busy, or a previous tape needs to be rewound and ejected, access times can stretch to minutes until a drive becomes available. If a tape drive is available, average access times to beginning of file can be a reasonable 30 to 90 seconds.

LTO-8 tape data transfer rates of up to 360 MB/s are quite high and superior to the practical bandwidth of many disk drives and gigabit Ethernet. Note that a high transfer rate is only achieved while the tape is able to stream. If read requests are not sequential, the tape drive has to stop streaming and seek to a new location on the tape and net bandwidth can drop significantly. Intelligent writing and reading of LTFS tapes by applications will be critical for performance-oriented solutions. Tape drive hardware-based data compression if appropriate for the content type(s) can increase the net data rate and capacity up to 250%.

The LTFS software provides information about the location of files on tape, enabling storage managers and other applications to properly sequence read requests in 'tape order'. During file writes, it is important for future read performance that files are written contiguously. If multiple users or applications are writing to the same LTFS tape at the same time, the blocks of their files will interleave on the tape precluding the capability to do streaming reads. LTFS-savvy backup and archive storage management applications will manage multi-user requests to minimize contention and optimize performance.

Deleted or "overwritten" files remain on the tape since new data is always appended. On the positive side, this enables older versions of files to be recovered in case of inadvertent deletion; however, it also enables older files to be recovered despite intentional deletion. Deleted files can be truly deleted and their space can be reclaimed if remaining valid tape contents are copied to another tape and the entire tape is reformatted. These are general tape processes and not specific to LTFS.

As LTFS formatted LTO tape cartridge capacities have increased from approximately 1.425 TBs with LTO-5 to 11.4 TBs with LTO-8, the number of files that can be stored on the single cartridge as grow as well. An LTFS formatted LTO-8 could contain nearly 950,000 12 MB 2K DPX files. This could have performance implications working with tapes with very large numbers of files.

The LTFS Index Partition is processed at the time the tape is mounted affecting the initial availability of a tape's contents and again when the tape is unmounted. Periodically the Index is also written into the Data Partition affecting performance and capacity utilization. For use cases with very large numbers of relatively small files the use of containerization, if appropriate, could be beneficial, reducing the number of files LTFS is indexing and tracking.

To address this issue the LTFS 2.5 Specification includes support for Incremental Indexes. Implementation of this capability could reduce the performance and capacity impacts of the periodic writes of the Index Partition into the Data Partition.

LTFS Implementation Options

General LTFS Applications

Users have an ever-increasing number of options for applying the benefits of LTFS to their workflows. For the software developer or adventurer, one could write a DIY (do-it-yourself) LTFS application using the LTFS specification. This is not for the faint-of-heart and use of an independent LTFS format verifier would be highly recommended.

LTFS STAND-ALONE DRIVE SOFTWARE

For many sole practitioners or small organizations, LTFS Stand-Alone Drive Software with one or two LTO stand-alone tape drives maybe all they will require. Most of the LTO tape drive provider companies (e.g. HPE, IBM, Overland-Tandberg and Quantum) offer LTFS software drivers. LTFS software drivers enable stand-alone LTO drives to be accessed via the operating system. LTFS presents through the standard file system interfaces of the platform on which it runs (POSIX in Linux and macOS, Windows file system interfaces in Windows). As such, it can be used with Windows Explorer or Mac Finder or command-line tools and scripting environments that work with the native file systems. Vendors may provide free of charge downloads of the software drivers that typically support Linux, macOS and Windows environments. Some have included GUI shells and basic utilities.

For users wanting more digital media-aware solutions, a number of vendors have incorporated LTFS support for individual LTO tape drives into their products. The table following this section includes representative examples.

LTFS LIBRARY SYSTEM SUPPORT

Most medium to large scale users will want to take advantage of vendor supplied and supported LTO LTFS library implementations. Several LTO tape library providers have LTFS support for their LTO library systems, including HPE's StoreOpen and IBM's Spectrum Archive, Quantum's Scalar LTFS Appliance, Spectra Logic's Black Pearl appliance, and XenData. These offerings combine the tape content access of the single drive LTFS versions with control of the tape library. All LTFS formatted cartridges in a tape library are visible and accessible. Other vendors provide solutions utilize LTFS to present LTO libraries as Network Attached Storage (NAS) systems.

There are also digital media-aware solutions with integrated LTFS support for LTO library systems. The more advanced solutions in this category include support for different storage types (Flash, disk, tape and cloud), multi-library, and multi-site functionality. This allows an enterprise to implement its own best practices digital asset retention and archiving solution as described early in this paper. The table following this section also includes examples of these providers.

LTFS Products for Media and Entertainment

Many solution providers support LTFS. Traditionally, data tape was only appropriate for archive use cases, but LTFS gave LTO tape advantages for broader use. This section provides a representative list of products and services focused on the media and entertainment market with LTFS implementations. The list is as of August, 2019, and due to the dynamic marketplace is necessarily incomplete. The applications and services are listed alphabetically within categories. (Additional LTFS implementations beyond media applications can be found at <https://www.lto.org/technology/ltfs/ltfs-Implementers/>)

COMPANY	PRODUCT(S) / SERVICES
1 Beyond, Inc. www.1beyond.com	Wranglers NetDrive ThunderTape, ThunderMax also EzStor (archive solution)
Atempo www.atempo.com	Atempo ADA
Echoleaf echoleafsystems.com	Topan Studio Forever Archive Gemini Netezza Archive
FOR-A www.for-a.com	LTO Server – LTO-80 LTO Server – LTS-70
HPE HPE-StoreOpen	HPE StoreOpen and LTFS software
IBM www.ibm.com	IBM Spectrum Archive Single Drive Edition IBM Spectrum Archive IBM Archive and Essence Manager (AREMA)
Imagine Products www.imagineproducts.com	PreRoll Post myLTO
Masstech www.masstech.com	FLASHNET MASSTORE
Oracle www.oracle.com/tape-storage/ltfs-faq	Oracle's StorageTek Linear Tape File System, Library Edition and StorageTek Linear Tape File System, Open Edition.
ProMax www.promax.com	Platform
Qstar www.qstar.com	Archive Storage Manager LTFS as NAS
Qualstar www.qualstar.com	Q1™, LTO Based Single Drive LTFS Archive Appliance
Quantum www.quantum.com	StorNext Advanced Data Management
Spectra Logic spectrallogic.com/products/blackpearl	Black Pearl Converged Storage System front-ending LTO/LTFS tape libraries.
StorageDNA www.storagedna.com	DNAevolution (Data management solution for the file-based media workflow.)
StrongBox Data Solutions www.strongboxdata.com	StrongBox (A cache-based LTO/LTFS data archive solution that presents to users as a NAS.)

Tolis www.tolisgroup.com	BRU Producer's Edition (Manages creative artist session archives.)
TransMedia Dynamics www.tmd.tv	Mediaflex Paragon - Digital and Media Asset Management and Workflow Solutions
XenData xendata.com	XenData SX Series archive servers XenData6 Workstation software
YoYotta www.yoyotta.com	YoYotta LTFS YoYotta Automation
SERVICE PROVIDERS	IMPLEMENTING WITH LTFS
Arkivum www.arkivum.com	Arkivum Assured Archiving (Providing a media-oriented archive cloud service.)
Digital Bedrock www.digitalbedrock.com	Digital Preservation Service

What Might the Future Hold for LTFS?

- LTFS will continue as a principal format for media data interchange and long-term digital media archiving.
- LTFS enabled solutions for digital media workflows, backup, and archive storage management applications will continue to be enhanced for:
 - new LTO generation support
 - the addition of new LTFS Specification features.
- Solutions integrating LTO/LTFS-based libraries with Flash storage and AI could provide low-latency, cost effective work-in-process active archive repositories.

References

2011 TECHNICAL EMMY AWARD

<http://www.emmys.com/articles/video-highlights-engineering-emmys-host-mayim-bialik>

<http://www.prnewswire.com/news-releases/ibm-and-fox-networks-group-win-engineering-emmy-award-for-broadcast-technology-innovation-132332528.html>

ADDITIONAL INFORMATION

LTO Program
www.lto.org

SNIA "Linear Tape File System (LTFS) Format Specification"
https://www.snia.org/tech_activities/standards/curr_standards/lvfs

“The Linear Tape File System”, David Pease et al.
www.researchgate.net/publication/224151009_The_Linear_Tape_File_System

“Long Term Data Retention – Key Considerations”, Quantum Technology Brief
<https://www.mercurymagazines.com/pdf/QUANTUM1.pdf>

“In Search of the Long-Term Archiving Solution – Tape Delivers Significant TCO Advantage over Disk”,
Clipper Group
<http://www.clipper.com/research/TCG2010054.pdf>

“File based media workflows using LTFS tapes”, Arnon Amir et al.
Association for Computing Machinery (ACM)
Proceedings of the International Conference on MultiMedia 2010
<http://dl.acm.org/citation.cfm?doid=1873951.1874269>

“Using Tape to Optimize Data Protection Costs and Mitigate the Risk of Ransomware for Data-Centric
Organizations”, Phil Goodwin, IDC; April 2018
<https://www.lto.org/wp-content/uploads/2018/04/US43710518.pdf>

“Quantifying the Economic Benefits of LTO-8 Technology”, Alex Arcilla, Enterprise Strategy Group, August,
2018

Wikipedia entry
http://en.wikipedia.org/wiki/Linear_Tape_File_System

ABOUT THE AUTHORS

George Anderson and David L. Trumbo are both principals of Media Technology Market Partners LLC,
doing business as MTMP Solutions.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the contributions of the primary author of the original study, Rainer Richter and of various members of the LTFS community especially Dr. David Pease. Rainer was a principal in MTMP Solutions leading the company’s software development and integration services. Dr. Pease, retired IBM Distinguished Engineer in Exploratory Storage at IBM’s Almaden Research Center, led the team that developed the initial LTFS specification and prototypes. While at IBM he continued to play a significant role in the SNIA LTFS Technical Working Group.

ABOUT MTMP SOLUTIONS

MTMP Solutions provides system design, integrated solutions, consulting services, hardware, and software to the media and entertainment industry. The firm specializes in storage solutions for digital content with an emphasis on LTFS-enabled solutions. MTMP Solutions was an early advocate and is a continuing evangelist for the use of LTFS in storage solutions not only for media and entertainment but all digital media content. MTMP Solutions was founded in 2009 by George Anderson and David L. Trumbo. The firm is headquartered in Los Angeles, California, www.mtmpsolutions.com.

Appendix

Storage Tiers

Tier	Latency	Bandwidth	Cost / GB	Typical Technology
Tier 1 (Online)	Very low (milliseconds)	Very High	High	NVMe, Flash, SAN, High Speed NAS
Tier 2 (Nearline)	Low (seconds)	High	Medium	JBOD, NAS
Tier 3 (Archive/ Offline)	High (minutes)	High	Low	Tape, Cloud Archive

While tape has always been a natural choice for the offline tier, LTFS enables tape to be a strong candidate for the nearline tier when integrated in a networked appliance. LTFS enables applications developed for the traditional nearline disk tier to access data on LTFS tape without being aware the data is actually on tape. Latency (but not bandwidth) is the key performance metric which determines suitability for a particular use case.

LTFS Success Stories:

Here are three success LTFS implementations in digital media applications. For more LTO success stories visit <https://www.lto.org/resources/case-studies/>.

Fremantle N.A. – America’s Got Talent

COMPANY PROFILE

With production units throughout their global offices, Fremantle is one of the biggest creators, producers and distributors of television programming in the world. The North American production unit based in Burbank, California, produces scripted, non-scripted and game shows including ‘American Gods’, ‘American Idol’, ‘America’s Got Talent’ and ‘The Price Is Right’ just to name a few.

BUSINESS NEEDS

- For ‘America’s Got Talent’ (AGT), now in its 14th season, Fremantle required a cost-effective way to store its large video archive on near-line storage.
- An open standard method for archiving and protecting all original camera and post-production mezzanine files for the current and past seasons
- Straightforward access to all those files for the show’s post production team of editors and assistant editors.
- The entire solution had to reside in a single rolling rack to support remote production locations as well as in-house operations.

SOLUTION - RESULTS

- AGT implemented a NAS-LTFS appliance with a (48) slot LTO library. The solution is capable of multiple petabytes of storage between tapes in the library and tapes 'on the shelf'.
- The system presents all content to the editorial team as network shares like any other network attached storage (NAS). All the LTO tapes are formatted with LTFS.
- The solution includes an expanded disk cache to lower latency access to current season content.
- LTO tape storage provides significant cost savings, increasingly as the archive grows with each new season's content.
- Two copies are automatically made on separate tapes. One copy is stored offsite for protection from natural disasters.
- Second copy tapes are stored offline to protect content against hackers, ransom-ware and malware attacks.
- Entire solution – appliance, LTO library, expanded disk cache and (2) UPSs (Uninterruptable Power Supplies) – fits in a single rolling rack. It has already moved 5 times mid-season to support live show production.
- To date the solution is managing over 3.1 million files requiring 1.63 PBs of LTO/LTFS tape storage.

Silverback Films

(original published in June, 2019 at <https://www.lto.org/resources/case-studies/>)

COMPANY PROFILE

Silverback Films specializes in the production of high quality wildlife films for both television and cinema. Formed in 2012 with productions including "Disney Nature Penguins" and the "Our Planet" series, Silverback Films brings together a world class team of wildlife film makers to create high quality history films.

BUSINESS NEEDS

- Silverback Films required an industry standard storage format to deliver content to clients
- Needed to lower storage costs—not cost effective to keep large video archive in on-line storage
- Desired easy process to retrieve digital content from video to archive
- Wanted secure method to protect valuable video assets from disaster and sabotage

SOLUTION - RESULTS

- Silverback Films installed two LTO tape automation libraries with an LTFS enabled system capable of multiple petabytes of storage
- System allows proxy file editing that conform back to raw footage on LTO tapes formatted with LTFS
- LTO - LTFS tapes provide industry standard format for content delivery to clients
- LTO tape system provided significant storage cost savings
- Easy to make two copies of tapes with one offsite to protect assets from natural disasters
- Tapes are offline protecting content against hacker and malware attacks

Digital Bedrock

COMPANY PROFILE

Digital Bedrock is a managed digital preservation service provider, preserving high-value digital content for studios, media and entertainment producers, museums, distributors, and archives. Founded in 2016, it has securely preserved millions of files for Academy Award-winning films and Emmy Award-winning episodics, including pre-release content. Preservation actions include monitoring bit health and format obsolescence vulnerabilities over time, and migration to new storage media in the future.

BUSINESS NEEDS

- Digital Bedrock needed a stable storage medium for storing clients' files after the preservation process.
- Target storage needed to be easily integrated into Digital Bedrock's developed digital preservation software platform.
- Security be strong; clients' assets cannot be accessed by anyone other than by Digital Bedrock staff, not even by the client.
- In the Digital Bedrock system, three or more verified copies are created and stored in geographically dispersed locations for disaster recovery; storage costs must be low, and physical locations must be identified.
- SHA-512 checksums on the stored files must be verified annually ("fixity checks") to substantiate data integrity over time.
- Assets must be easily accessible without a long delivery time to fulfill client requests.
- Simple "Exit" and storage migration strategy are needed, without a storage vendor lock-in.

SOLUTION - RESULTS

- Digital Bedrock installed two LTO tape libraries for writing three copies simultaneously, as well as reading files off clients' delivered tapes and performing annual fixity checks.
- The company built its own LTFS tool to control all tape library drives for writing, reading, and performing fixity checks over time. The tool is integrated with the company's digital preservation platform.
- Using LTFS enabled storing the files in exactly the directory structure as the client provided, making future migration easy.
- Storing client data offline on LTO media and in locked, secure facilities means the data cannot be hacked. The distributed physical storage locations and personnel with access are known.
- Accessing files from the LTO tapes is faster than the process to use cloud "cold" storage, which requires first moving the files from cold to standard storage, then downloading them.
- LTO media is stable; no files have failed their fixity check after verifying millions of files.

About MTMP Solutions

MTMP Solutions provides system design, integrated solutions, consulting services, hardware and software to the media and entertainment industry. The firm specializes in storage solutions for digital content with an emphasis on LTFS (Linear Tape File System) enabled solutions. MTMP Solutions was an early advocate and is a continuing evangelist for the use of LTFS in storage solutions not only for media and entertainment but all digital media content. MTMP Solutions was founded in 2009 by George Anderson and David L. Trumbo. The firm is headquartered in Los Angeles, California.

For more information, visit www.mtmpsolutions.com

T: +818.763.6478 (U.S. & Canada)

+1.818.763.6478

E: contactmtmp@mtmpsolutions.com

