

Facts About LTO and LTFS

Why LTO and LTFS may be the best option for long-term storage

Many media professionals face a number of challenges as a result of the phenomenal growth of file-based workflows and media. The need to reliably and cost-effectively store content has become critical for production and post production facilities transitioning to file-based pipelines.

Challenges of moving to a file-based pipeline

- + Most users have to store growing amounts of file-based media, but keep storage costs down
- + Digital assets may have future monetary or historic value and need to be protected long-term
- + An archive and retrieval solution is required that enables users to quickly find and restore media when needed

Storage options

In addition to Linear Tape-Open (LTO) and Linear Tape File System (LTFS), here are the other existing storage options for the purpose of comparison:

- + **SAN or nearline disk arrays:** Many media professionals leave their content on a SAN, or use lower cost nearline disk arrays. While this may be a short-term solution, it becomes expensive in the long run due to the costs of storage, scaling, and power-on. SAN and nearline disk systems can cost upwards of \$1.50/GB. Compare that to the alternatives of removable hard drives and LTO which cost approximately \$.02/GB – \$.03/GB.
- + **Removable hard drives:** Another option adopted by a number of facilities moving to file-based pipelines is removable hard drives. Hard drives are inexpensive and easy to work with, however; removable drives are not designed or built for long-term archival of content. This is due to the inherent architecture of hard drives that makes them prone to failure due to drops or inactivity on a shelf.
- + **Optical:** While optical storage is another option, it offers poor density (50 GB) and higher costs when compared to LTO or removable drives. As a result, it is often better to pick an alternative to optical for long term file-based archiving needs.

Introduction

LTO-5 and LTFS

Linear Tape-Open (LTO) is a standards consortium that dictates the specification for digital data tape. LTO is tape built for storing files rather than analog data. LTO drives and tapes are standardized and manufactured by a number of IT vendors. Tape drives are manufactured by HP®, IBM®, and Quantum®, and vendors such as Imation®, Fuji®, and Maxell® build LTO tapes. LTO has been standardized for IT workflows for the past decade with millions of tapes in use. LTO-5 is the fifth generation that can store 1.5 TB of raw data per tape. Detailed specifications for LTO-5 can be found at: http://en.wikipedia.org/wiki/Linear_Tape-Open.

LTFS: a file system for tape

Linear Tape File System (LTFS) is a recent development from the LTO consortium. It is designed to enable an open and more efficient way to access tape. Here is an example to put this file system into perspective: many years ago, hard disk-based systems standardized on file systems. These file systems include DOS, NTFS on Windows®, HFS on Mac® OSX, and XFS on Linux®. A file system defines a format for how content should be written to disk. A file system also provides file system driver software that “mounts” the file system, and makes it available for applications to read and write files. These applications can be utility-based like Finder®, and Explorer that enable you to more easily copy, move, and rename files. Additionally, file systems allow any application to modify, and access files. File systems have enabled users to easily access files on hard drives, and across one or more workstations.

Until now, a standard for tape-based mediums did not exist. LTFS is the first standard that defines a file system for tape. Much like disk-based systems, LTFS defines an open format for how it writes content to tape. LTFS also consists of file system driver software that allows tapes to be “mounted” and be seen like any other mounted device.

More details can be found at: http://en.wikipedia.org/wiki/Linear_Tape_File_System.

Benefits

LTO and LTFS for long-term archiving

A perfect storage medium does not exist. Compared to the other available solutions, LTO offers the most reliable, cost-effective alternative to solve the challenges of long-term data management for file-based workflows.

LTO and LTFS for media professionals

LTFS enables LTO to look more like disk, and this allows LTO to act not only as an archive medium to protect assets long-term, but also as a nearline medium for storing raw and post assets. With growing amounts of file-based content, media professionals need a cost-effective, reliable and accessible storage medium. LTFS combined with LTO tape makes this possible.

Cost of LTO-5

LTO-5 data tapes cost approximately \$55.00 USD per 1.5 TB of tape. This means that LTO-5 costs roughly \$.03/GB – \$.04/GB. That is roughly 30 times less than the least expensive disk array system and about the same cost as removable hard drives. This allows you to significantly reduce the cost of storing large amounts of file-based media and keep this cost low as you scale.

Reliability of LTO-5

Unlike a removable hard drive that is a complex device with heads and electronics, an LTO tape is a simple spool of tape with no moving parts. The read/write technology is hosted in a tape drive and not in the cartridge (like a hard drive). This means that an LTO tape has a far greater chance of surviving a disaster than a hard drive on a shelf. Further, LTO-5 magnetic tape is designed to last on a shelf for 30 years (provided the appropriate temperate and humidity controls are maintained). As a result, media archived on LTO has a far greater chance of survival than a hard drive.

Long-term, open access via LTFS

Prior to LTFS, a major drawback cited against LTO was the difficulty of access and closed data format on tape. As long as you write data to LTO-5 using LTFS, your data stays open as files, similar to a hard drive. LTFS has made LTO as open as any removable hard drive.

Performance of LTO-5 tapes

An LTO-5 tape drive can read and write data at an impressive 1.2 Gib/s, (approximately two times the speed of a hard drive). This means that an appropriately configured system and software can achieve the following archive and restore times for different media file types:

Archive and/or restore times for 1 hour of:

- | | |
|--------------------------------------|--|
| + Avid® DNxHD 145: 9 minutes | + DVCPro HD (Supported by a number of camera manufacturers): 8 minutes |
| + Avid® DNxHD 220: 12.5 minutes | + ProRes HQ (Apple® codec): 16 minutes |
| + Panasonic® AVCIntra 50: 4 minutes | + ProRes 4444 (Apple® codec): 21 minutes |
| + Panasonic® AVCIntra 100: 8 minutes | + REDCODE-36 (RED® camera codec): 18 minutes |

Built-in check-summing

LTO-5 is designed for archiving media long-term and for performance. To ensure that media is protected, LTO employs verify-after-write technology: a read head that follows the write head, and verifies the data just written. This provides an automated check-summing operation for all data written to LTO. In addition, LTO-5 provides twice the read/write performance of removable hard drives. Unlike a hard drive, where all data has to be read back to verify that data has been correctly written, LTO provides this capability during an archive, saving an entire read and restore operation. This means that, on average, an archive and checksum operation with LTO will be four times faster than performing an archive and checksum to a removable hard drive.

Store multiple hours of content on a single tape

LTO-5 has a raw density of 1.5 TB. This means that an LTO-5 tape can hold multiple XDCAM discs, P2 cards and hours of traditional video tape footage. This is critical in reducing storage costs. The typical number of hours of footage you can fit on a single LTO-5 tape is as follows, (assumption: 1.2 TB useable per LTO tape; rest is factored in for LTFS headers and other metadata).

1 Lto-5 can hold:

- + 95 hours of AVCHD
- + 49.95 hours of AVCIntra50
- + 26 hours of AVCIntra100
- + 25 hours of DVCProHD
- + 61 hours of XDCAMHD
- + 13.5 hours of ProResHQ
- + 22 hours of DNxHD 145
- + 9 hours of ProRes 4444
- + 10.5 hours of REDCODE-36

Myths and drawbacks

LTO tape drives read back two generations of LTO tape

Every generation of LTO can read back two previous generations of LTO tapes. An LTO-5 tape drive can read back LTO-3 and LTO-4 tapes. The tapes are specified for a 30-year shelf life, and as a result, you have to ensure that you also keep the tape drives around for 30 years.

This is often cited as a drawback of LTO technology compared to traditional film or video tape. However, it is considered good practice to retransfer your media every five to ten years to truly guarantee the long-term health of your media. Instead of having to face years of untouched analog tape media, LTO technology ensures we check and refresh storage media every two generations.

Additionally, LTO enables you to take advantage of cost and performance efficiencies every generation. LTO doubles its capacity and speed every two to three years enabling you to store your media in half the space.

Can't use LTFS directly, like a hard drive

A misconception regarding LTFS technology is that it can be used like a hard drive. LTFS makes a tape look like a hard drive, but this is misleading. These are vastly different storage mediums: disk is random access and tape is not. Every application that writes to a file-system expects that a disk, (a random access device), is behind the file-system. Without the knowledge that there is a tape behind a file-system, an application will experience a very poor I/O response. This is the reason why many end users deemed that LTFS is not ready for prime time. They pointed their favorite application directly, (such as Finder or Explorer), and as a result, simple list, read, and write simply hung.

LTFS is a major breakthrough in that it will herald in a new set of applications for tape, but it is not a drop-in replacement for disk.

LTFS versus tar

Like tar, LTFS is an open standard and format for tapes. Therefore, if one vendor wrote to a tape using tar or LTFS, another user could independently extract the content. The most significant difference between LTFS and tar is that LTFS makes tape more accessible.

While this has benefits for open access, this feature of LTFS alone does not offer a compelling advantage over tar.

The true benefits of LTFS come from its ability to empower application vendors to build better workflows, (powered by LTO), for the end user community. As better workflows are built on top of LTFS, end users are likely to prefer LTFS over tar-based pipelines. Early adopters of LTFS may have the advantage of being aligned with the future direction of the industry.

Conclusion

As file-based acquisition becomes the norm, it is essential for media organizations to develop a long-term content archival and retrieval strategy. Today, LTO and LTFS is a cost-effective and reliable option to protect digital media assets. This proven technology will advance, the performance will accelerate, and it will continue to be a powerful storage solution for the film, video and broadcast industry.