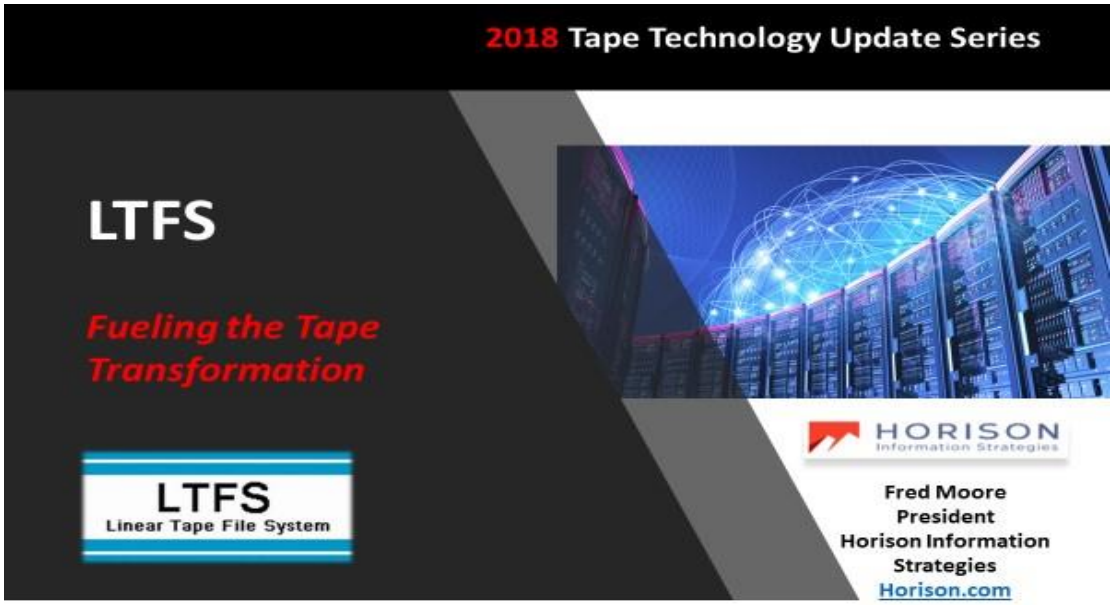

Abstract

Since its inception, tape has been a sequential (linear) access storage device. Though this proved beneficial for backup and many "high performance data streaming applications", the ability to randomly navigate a tape cartridge to locate a specific file or object was inefficient. The LTFS™ (Linear Tape File System) Single Drive Edition (LTFS SDE) format was developed by IBM and adopted by the LTO Consortium (IBM, HPE, and Quantum) to address this challenge. On April 12, 2010, LTFS was announced at the National Association of Broadcasters (NAB) annual trade show for the LTO-5 standard tape format. The LTFS format specification is fully supported by the LTO roadmap which defines the specifications for the entire future LTO tape drive family. LTFS now provides a standard open file system for LTO drives, the latest enterprise TS11xx and T10000 drives from IBM and Oracle StorageTek, and tape libraries from several storage suppliers, further validating and expanding its value. IBM continues to be the principal developer of LTFS and offers LTFS as open source software in a variety of editions for single drive and robotic tape libraries. IBM officially renamed their LTFS solution as IBM Spectrum Archive™ on May 5, 2015. Other tape system suppliers offer their own editions of LTFS that are compatible with IBM's version.

***Key point:* LTFS provides the tape industry with its first standard open systems file system providing easy and faster access to LTO and enterprise tape files.**

How Does LTFS Work?

LTFS was not the first tape file system. Prior to LTFS, Tar (Tape Archive System) was the most commonly used tape file system. Tar doesn't provide metadata indexing capability like LTFS. Identifying cartridge contents using Tar can take several minutes if not hours compared to an LTFS cartridge, which takes a few seconds. With Tar, the user must scan the entire written area of tape to discover what data resides on that tape. The files in a Tar archive are not compressed, just gathered together in one file called a "tarball" and was mainly used to transfer files among UNIX systems. Untangling a tarball can take time!

LTFS provides significant improvement over its predecessor(s). LTO tape cartridges are physically laid out with four, wide data bands sandwiched between five narrow servo bands for alignment. The tape head assembly that reads from and writes to the tape straddles a single data band and the two adjacent servo bands for optimal alignment. The tape head can have 8, 16, or 32 data read/write head elements and 2 servo read elements. The set of 8, 16, or 32 tracks are read or written in a single, one-way, end-to-end pass that is called a "wrap". The tape head shifts laterally across the tape to access the different wraps within each band and to access the other bands.

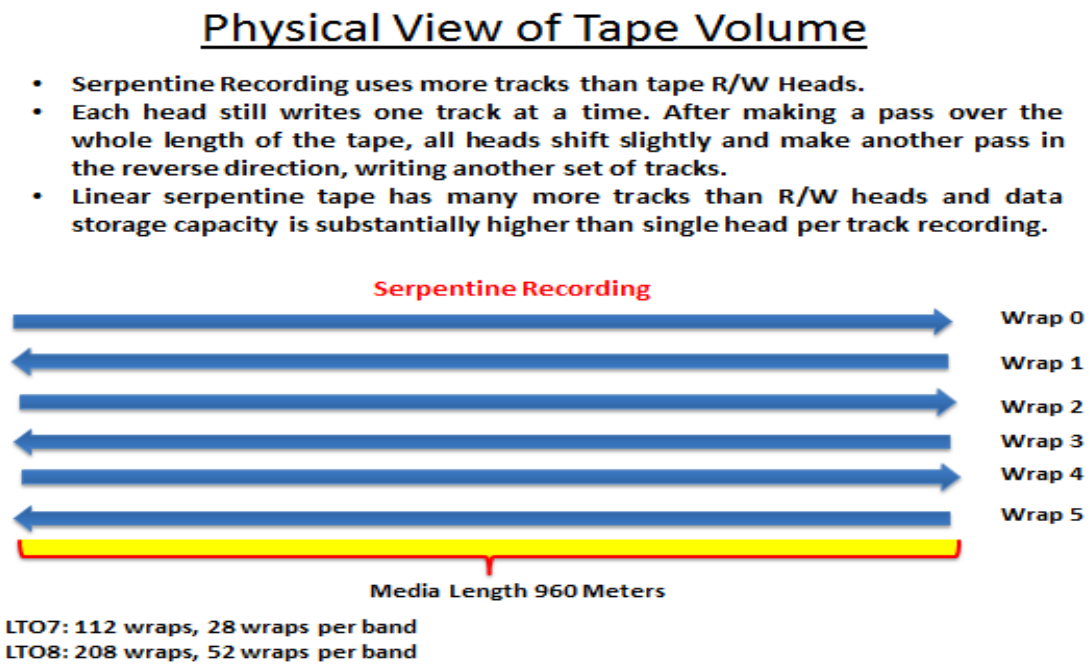

LTFS creates two partitions on serpentine tape, one partition holds the content and the other(s) holds the content's index. A specific application is no longer needed to determine what's written on the tape because the metadata (the descriptive information of what is on the tape) in the Index Partition describes the files, data contents and pointers to their physical tape location in the Data Partition. The LTFS metadata in the Index Partition enables faster searching and accessing the files in the second partition via a GUI (Graphical User Interface) provided by the LTFS operating system. An LTO tape cartridge enables the creation metadata or "tagging" of individual files providing rapid file access making it possible to double-click, open, drag, drop, copy and paste files to tape in the same familiar way as performed on disk, SSD, CDs, or DVDs.

The two LTFS partitions are shown in the chart below. Note: LTO-6 extended the specification to allow 4 separate partitions. Physically, the Index Partition is physically organized across the full length of the LTO tape media and requires about 5% of the total cartridge capacity. The Index Partition is created to use one whole wrap on the media. The LTFS partitioning scheme uses an additional wrap as a Guard Wrap between the Index Partition and the Data Partition. The Guard Wrap serves as a physical buffer zone to ensure that writes to one partition cannot interfere with data on the other partition.

Logical View of LTFS Volume

- LTFS utilizes media partitioning (LTO-5+, the T10000, and TS11xx family)
- The LTFS tape is logically divided into partitions "lengthwise".
- LTFS places the index in first partition and file(s) in the second partition.
- The LTFS index enables faster searching and accessing the files in the second partition via a GUI (Graphical User Interface).

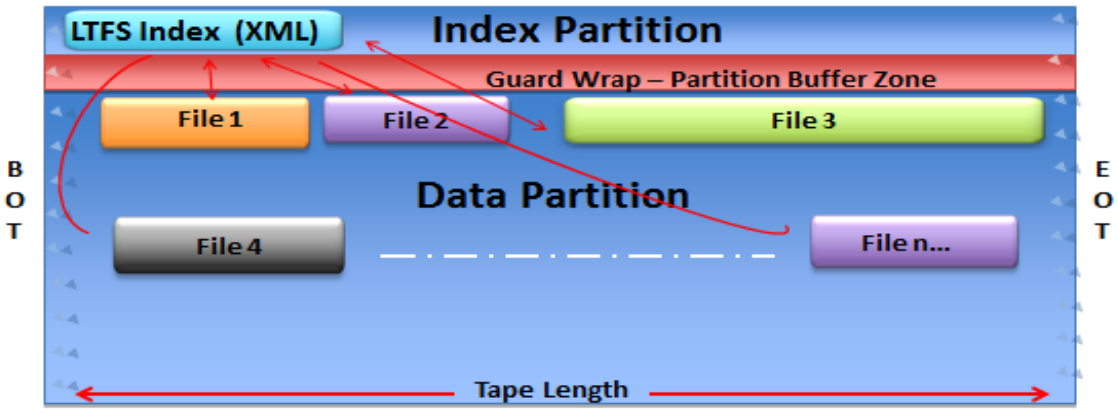

When a tape with an LTFS index is inserted, the LTFS will load the metadata for all the active files on the tape cartridge into server memory for faster future access. When a new tape cartridge is mounted, or the data on an existing physical tape is updated or changed, the LTFS index in server memory is updated with that cartridge's information first, then asynchronously that index is backed up and written to the tape cartridge for consistency and higher availability. LTFS will recognize when tape leaves and reenters the library and performs consistency checks to determine if tape index has changed.

LTFS Functionality Summary	
LTFS Drive **LTFS Library** **LTFS Enterprise** **LTFS Open Systems**	Ann. April 12, 2010 for LTO-5 drives (LTFS Single Drive Edition LTFS SDE) Ann. July 12, 2011 (LTFS Library Edition LTFS LE) Ann. June 10, 2013 (LTFS Enterprise Edition LTFS EE) Ann. Sept. 19, 2014 (LTFS Library Edition for Windows using third party libraries with LTO-5 and all subsequent LTO drives
LTFS Software Cost	Contact the specific LTFS vendor for any applicable pricing details.
Operating Systems	Linux, Apple OS X (Mac), Microsoft Windows.
Tape Drive and Library Support	Beginning with LTO-5 and all subsequent LTO versions, IBM TS1140, TS1150, TS1155, Oracle T10000C and T10000D. Most modern tape libraries support LTFS.
Tape Mgmt. SW	No additional tape management, backup software or utilities are required.
Functionality	Functions such as File Open, Write, Read, Append, Delete and Close from an application are supported on LTFS enabled tape.

Where to get LTFS	Download LTFS from the vendor's websites - (IBM, HPE, Oracle, Quantum). LTFS versions are available as open source and can be downloaded from the following vendor URLs: For HPE StoreOpen and Linear Tape File System (LTFS) Software For IBM Spectrum Archive For Oracle/STK LTFS For Quantum LTFS
LTFS Support for Prior and Future LTO Generations	LTFS tape partitioning was introduced with the LTO format beginning with LTO-5, and therefore, earlier generations such as LTO-1 through LTO-4 drives do not support LTFS. Note: The LTO roadmap indicates that LTFS tape partitioning will be supported in all future LTO drives.
Cartridge Contents Inquiry for Fast Access and Retrieval	After the tape library is initially inventoried in memory, tape cartridges do not have to be remounted to retrieve its content index data and basic cartridge information like volume name, date, serial number and pointers saving time.
Partition Update Capability	Each partition (Index and Data) can be accessed and updated independently. The ability to access small sections of data on tape is significantly improved.
Protection for Index Partition	With LTFS Version 2, the index partition is periodically copied to the data partition on the tape for backup in case the primary copy is unavailable.
Data Sharing	Files stored on LTFS can be shared between multiple operating systems, different LTFS editions, and between different applications.

Source: Horison, INC.

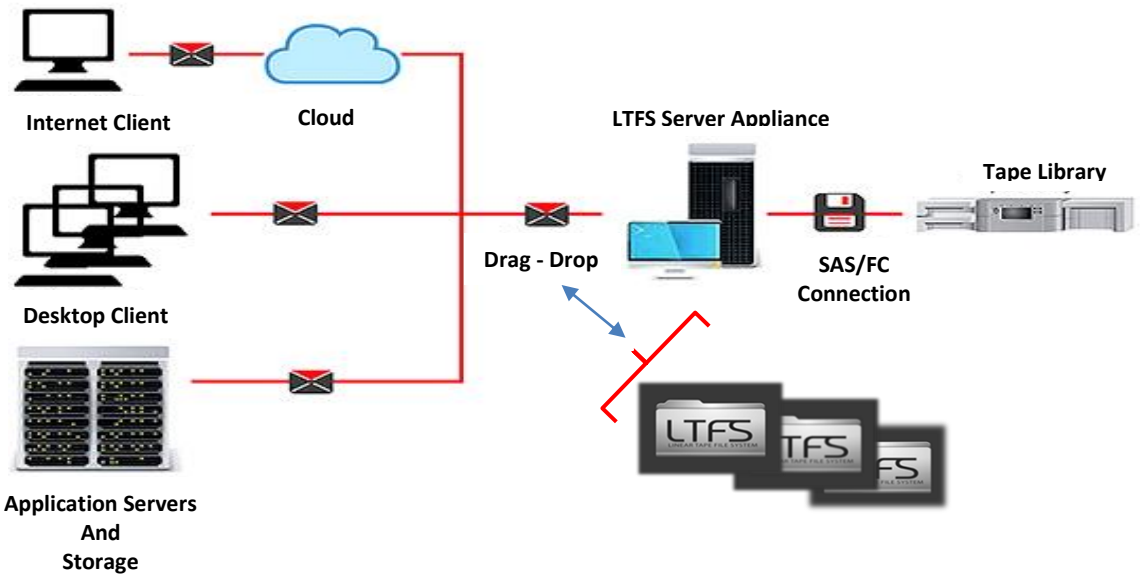

Key point: ***LTFS is <u>not a disk replacement</u> but it changes the rules for tape access and enables much faster tape access than any previous tape format. With LTFS, tape looks and acts more like disk.***

LTFS Addresses New Storage Intensive Applications

The decision to make LTFS available as open source at no cost, instead of waiting for software companies to add their support, has been quite effective while allowing ISVs to add support as required. Historically, tape usage has centered on backup, restore, disaster recovery, archiving and long-term data preservation applications.

As a result of LTFS being open source, several new use cases for tape are arriving. Storage intensive applications consisting of unstructured data, 3-D images, multi-media, video, social network content, Big Data, surveillance and compliance data are experiencing a CAGR (Compounded Annual Growth Rate) of >50% annually and are ideal candidates for modern tape. The enormous wave of Big Data applications and the reality of the IoT has arrived. According to IDC's Digital Universe Study, digital data created should reach 44 zettabytes ($44x10^{21}$) by 2020 and forecasts that by 2025, the entire global data-sphere (all data created) will grow to 163 zettabytes (one trillion gigabytes) though most of this data is transient and will not be permanently stored or result in storage demand.

LTFS Assist for the Cloud

Cloud applications using tape have significant growth potential with LTFS. Tape cartridges are commonly used for long-term archival data storage due to their 30+ year media life, and tapes utilizing LTFS can often replace more expensive disk in the cloud as storing archival data is becoming a major cloud service (AaaS -Archive as a Service). The recent collaboration between two SNIA (Storage Network Association) technical working groups - LTFS and Cloud - is working to use LTO Technology with LTFS to deliver a more efficient cloud data transfer process by building a standard for LTFS Bulk Data Transfer. Cloud providers will be able to use this standard to develop interoperable services to transfer of very large files directly to tape using the LTFS format making "tape in the cloud" a standard and more cost-effective cloud offering. The much higher tape drive data rates compared to disk will also provide a huge assist to cloud service providers when moving large amounts of data consuming high cost bandwidth.

Spectrum Scale (GPFS) LTFS Support Key for HPC Systems

LTFS EE (Enterprise Edition) enables the use of LTFS as a fully supported storage tier in IBM's General Parallel File System (GPFS™) which was renamed as IBM Spectrum Scale on February 17, 2015. GPFS was developed by IBM as a high-performance clustered file system providing concurrent high-speed file access to applications executing on multiple nodes of clusters and is used by many of the world's largest supercomputer data centers. GPFS data management, integrated information lifecycle tools coupled with LTFS can manage petabytes of data and billions of files on disk and tape.

***Key point:* LTFS implementations can be as simple as a single LTO or enterprise tape drive, small and large tape libraries, or as part of an integrated tiered storage system using flash and hard disk. In this environment, data can move seamlessly between tiered storage layers based on data groups and policies.**

LTFS Now Embraces Object Storage

Object storage enables IT managers to organize archival content with its associated metadata into containers to easily allow retention of archival and unstructured data. Objects can be tagged with unique, customizable metadata so customers can easily manage storage consumption, cost, and security separately for each workload. As a result, on July, 2017 IBM Spectrum Archive™ Enterprise Edition V1.2.4 (LTFS) announced support, in connection with OpenStack Swift, to enable movement of cold (archive) data from object storage to more economical tape and cloud storage for long-term retention. LTFS now provides a back-end connector for open source SwiftHLM (Swift High Latency Media), a high-latency storage back end that makes it much easier to perform bulk operations using tape within a Swift data ring. Cloud storage is currently the most common use case for object storage.

Key point: ***LTFS support is gaining momentum and today there are 35 companies which are LTFS implementers. Expect the number of ISVs and cloud providers supporting LTFS in the future to increase.***

New LTFS-enabled Hybrid Solutions and Active Archive Gain Traction

LTFS is expanding its functionality into software-defined appliances coupled with data management software modules. An Active Archive is a hybrid solution which combines NAS, or disk arrays as a front-end cache, with tape and are most often combined with LTFS. An Active Archive makes disk and tape storage systems complimentary and can be viewed as tiered storage enabling access to data across a virtualized file system that automatically migrates data between SSDs, disk drives, tape, and the cloud. Active Archive solutions are available from several storage suppliers enabling a higher level of performance for more active archival files stored on tape.

Examples include:

- StrongBox Data Solutions
- Fujifilm Dternity
- HPE StoreEver Archive Manager Solutions
- IBM Spectrum Archive
- Oracle HSM
- Overland NEO Agility LTFS Archive Appliance
- Qstar Archive Manager
- Quantum Scalar LTFS and Artico
- Spectra Logic BlackPearl

To further address the enormous archive challenges, the Active Archive Alliance was launched on April 27, 2010 as a collaborative industry association formed to educate end user organizations on the new technologies that enable the most effective access to their archived data.

Key point: ***Hybrid solutions combining LTFS with disk, tape and SSDs have arrived providing faster online access, search capability and easier, more cost-effective retrieval of long-term archival data.***

The Tape Technology Renaissance is Well Underway

In addition to LTFS, the pace of tape technology innovation continues to be aggressive and the results are impressive. Tape recording demonstrations now indicate native capacities in excess of 200 TB per cartridge are achievable and project few limitations for areal density in the foreseeable future. Remember that tape scales capacity by adding media, disk drives scale capacity by adding more drives and the media life for modern tape media now exceeds 30 years.

The storage industry de-facto standard reliability measure of tape, BER (Bit Error Rate), has surpassed that of the best disk drive reliability by three orders of magnitude making both enterprise and LTO tape more reliable than the most reliable disk drive. Tape drive data rates are projected to be 5x greater than disk data rates by 2025. These and several other technological enhancements have given tape significant cost advantages and throughput advantages over disk and recent studies indicate as much as a 15x TCO advantage for tape over disk is common.

Key point: ***For the first time ever, tape is cheaper, more reliable, has much faster throughput, has a higher capacity and much longer media life than any of its disk counterparts.***

Conclusion

Throughout its history, tape has steadily evolved providing higher capacities and lower costs, but the enhancements were always built on a sequential access architecture limiting access time and constraining applications. LTFS is the first open file system that works exclusively in conjunction with tape technology and clearly sets a new standard for ease of use and portability for LTO and enterprise tape. Therefore, as future tape cartridge capacities continue to grow unbounded, the need for improved access and data retrieval capabilities using LTFS and the new advanced partitioning capabilities will also increase. Because of this and significant technological progress on many additional fronts in the past 10 years, the tape industry is effectively re-positioning itself to address many new high capacity, long-term storage and archive opportunities and promises to provide the most cost-effective storage solution for the foreseeable future.

Key point: ***Although LTFS is just beginning to scratch the surface of its potential, it can become one of the most significant developments in the tape industry.***

End of report