Abstract

Tape has been a sequential (linear) access storage device since its inception. Though sequential access proved beneficial for backup and many "high performance data streaming applications", the ability to randomly navigate a tape cartridge to quickly locate a specific file or object was inefficient. To address this challenge 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). On April 12, 2010, LTFS was announced at the National Association of Broadcasters (NAB) annual trade show for the LTO-5 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 is the industry standard open file system for LTO and for the latest enterprise TS11xx drives from IBM, 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 is the first open standard tape file system providing a faster and much easier way to navigate through LTO and enterprise tape data.**

What is a Tape File System?

A file system is a way of presenting data storage to an application in an intuitive and familiar way. Disks are formatted with a file system, making it possible to create files and folders with each file mapping onto the underlying physical blocks of storage. LTFS was not the first tape file system. Long before LTFS, the Tar (Tape Archive System) utility was first introduced with Version 7 Unix in January 1979 and was widely used. Tar doesn't provide metadata, tags or indexing capability like LTFS. Identifying cartridge

contents using Tar was complex and could take several minutes if not hours which compared to an LTFS cartridge, can be completed in seconds. With Tar, the user must scan the entire written area of tape to discover what files and data resides on that tape. This discovery process became untenable as cartridge capacities increased. The Tar archival files are not compressed, just gathered in one giant file called a "tarball" and were mainly used to transfer multiple files among UNIX systems. Untangling a tarball can take time!

LTO Media Format – Serpentine Recording, Bands and Wraps

LTO brings open standards along with many new and significant improvements to tape storage systems. LTO tape cartridges are *physically* formatted with four, wide data "bands" sandwiched between five narrow servo bands for alignment. The special dedicated timing-based servo patterns are written onto tape during manufacturing time and straddle the data bands. Servo bands serve like guard rails for the read/write head and provide compatibility and adjustment between different tape drives. 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 and track following. 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". Once a pass has been completed over the entire length of the tape, the heads all shift laterally across the tape to access either the different wraps within each band or to access the other bands in a "serpentine" (back and forth) manner.

Physical View of LTO Volume

- **Serpentine recording (back and forth) uses more tracks than tape R/W Heads.**
- **Each head still writes one track at a time.**
- **After making a pass over the whole length of the tape, all heads shift slightly and make another pass in the reverse direction, writing another set of tracks.**
- **Four data bands serve as guard rails for the read/write head providing compatibility and adjustment between different LTO tape drives.**

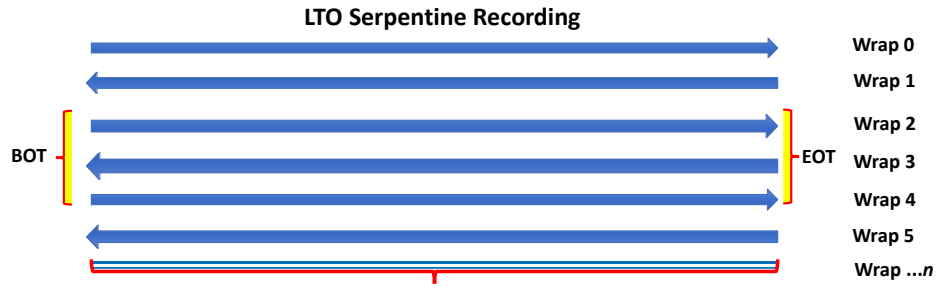

Gen.	Length (m, ft.)	Tracks	Tracks/wrap	Wraps	Wraps/band (4)	Tape Thickness
LTO-7	960/3149	3584	32	112	28	5.6 *u*m
LTO-8	960/3149	6656	32	208	52	5.6 *u*m
LTO-9	1035/3396	8960	32	280	70	5.2 *u*m

Source: Horison, Inc.

LTFS Media Format and Functions – Partitions, Metadata Tags, Drag and Drop

Using LTFS-capable media and drives, a tape can be mounted and treated as a file system and interacted with using files and folders, much like disk. This is achieved by formatting an LTO cartridge with LTFS and writing two or four separate *logical* sections on the tape, called "partitions". The first smaller Index

Partition, at the start of the tape, stores the list of filenames/folders, time stamps and a pointer to the blocks on the second, larger, data partition where the files are stored. When the tape is inserted (mounted) in the drive, the index information is quickly loaded from the first partition into the LTFS workstation or server memory. The LTFS operating system GUI (Graphical User Interface) displays the tape's entire contents enabling faster access to the requested files in the second partition. If 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 enables the creation of metadata or tagging of individual files in a cartridge. This enables faster file access to double-click, open, drag, drop, copy, and

Logical View of LTFS Volume

- **LTFS utilizes media partitioning (LTO-5+ and TS11xx Enterprise).**
- **The LTFS tape is logically divided into partitions "lengthwise".**
- **LTFS places the Index in first partition and Data file(s) in the second partition.**
- **The LTFS index enables faster search and file access in the second partition via a GUI.**

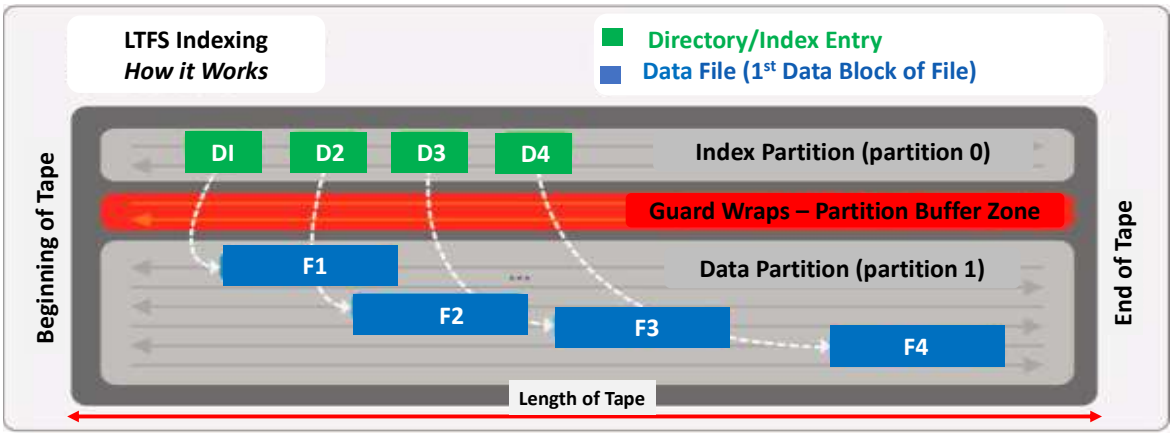

paste files to tape in the same familiar way as performed on disk, SSD, CDs, or DVDs via the LTFS GUI.

Two lengthwise LTFS partitions (Index and Data) are shown in the chart above. LTO-6 extended the specification to four separate data partitions running the length of the tape. The Index Partition spans the full tape length of the LTO tape media. An additional wrap, a Guard Wrap between the Index Partition and the Data Partition is used as a physical buffer safety zone ensuring writes to the index partition cannot interfere with files written on the data partition.

LTO and LTFS formatted media can be mixed within a tape library. LTFS recognizes when LTFS tape leaves and reenters the library and performs consistency checks to determine if the tape index has changed. While LTFS can make a tape appear to behave like a disk, the fundamental sequential nature of tape is unchanged. Unlike disk, new files are always appended to the end of the tape. If a file is modified and overwritten or deleted from the volume, the associated tape blocks used are not freed up: they are marked as unavailable, and the used volume capacity is not recovered. Data is deleted and capacity

physically recovered only if the whole tape is re-written. As a result of the new hardware and software tape functionality, an LTFS implementation is particularly suited to:

- Archival file storage using a simplified, self-describing format on sequential access media for applications requiring faster access times to tape files
- Transfer a large volume of data to or from a public, or private cloud
- Data export and import, data interchange and exchange between different operating systems
- Transfer a large volume of data to or from a remote location
- Direct file and partial file recall from sequential access media
- Update or synchronize a subset of data stored at a remote location

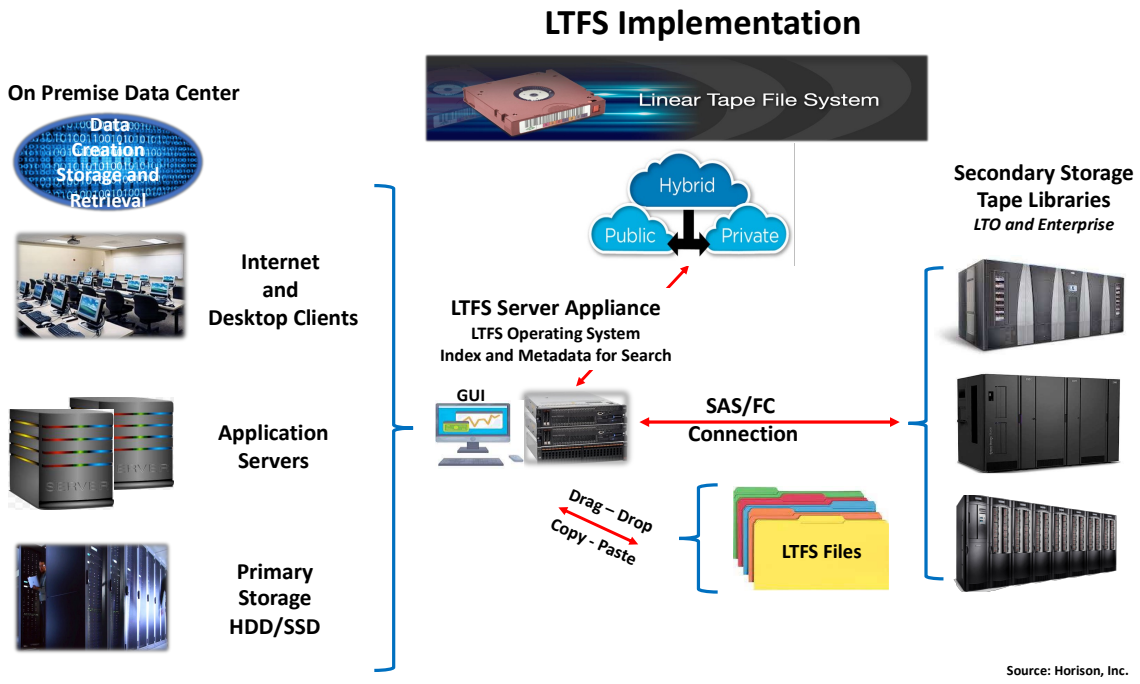

***Key point:* With LTFS, tape still requires sequential access to data but acts more like disk. LTFS is not a disk replacement, but it improves time to first byte compared to any previous tape format.**

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 allowing individual ISVs to add support when 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 have arrived. Any industry where vast amounts of data are being collected and stored can benefit from LTFS. For example, storage intensive applications including AI/ML, multi-media, video, 3D images, social media, surveillance, finance, compliance, Big Data, and archival storage for cloud service providers are experiencing a sustained CAGR of 30% or more annually and are ideal candidates for modern LTFS tape. Industry estimates vary somewhat but IDC's Data Age report estimates the total amount of digital data to be generated WW in 2025 to be ~175 ZBs with ~11.7 ZB actually stored on SSDs, disk and tape systems.

Additional industry estimates project that as much as 80% of the stored data (~9.3 ZB) is classified as archival "cold" data representing the immense secondary storage (tape) market, the largest and fastest growing storage segment. The role of LTFS powered tape in the zettabyte era adds significant value to cost-effectively manage the soaring secondary storage challenge.

Key point: ***LTFS has 34 companies which are licensed LTFS implementers. Expect the number of ISVs and cloud providers supporting LTFS in the future to increase.***

LTFS Manages Object Storage

Object storage is gaining momentum as an archive format and has become the de-facto standard format for storing data on tape and in the cloud. 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 to 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. The IoT generates much of its content in object format and the ability to archive and quickly retrieve this data for analysis from tape archives and data lakes provides a highly cost-effective use case.

LTFS Assist for the Cloud

LTFS can cost-effectively replace more expensive energy consuming disk farms in the cloud as storing archival data has become 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 plan to use this standard to develop interoperable services to transfer very large files directly to LTFS formatted media 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 additional benefits to cloud service providers when moving large amounts of data (ingress and egress) over expensive bandwidth connections.

Spectrum Scale (GPFS) Supports LTFS 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 HPC and supercomputer data centers. GPFS data management and integrated information lifecycle tools coupled with LTFS can manage petabytes of data and billions of files on disk and tape.

The Tape Technology Renaissance is Well Underway

In addition to the benefits of LTFS, tape technology innovation continues to be aggressive and impressive. Since the introduction of LTO-1 in 2000 with a native capacity of 100 GB, LTO capacity has

increased 180 times and data rates have increased by 20 times. Tape recording demonstrations from IBM and Fujifilm indicate native capacities up to 580 TB per cartridge are achievable. Remember tape can scale capacity simply by adding media, disk drives can *only* scale capacity by adding more drives and more energy consuming resources. LTO-9 and enterprise TS11xx data rates have reached 400 MB/sec. doubling that of most disk drives. The storage industry de-facto standard reliability measure, BER (Bit Error Rate), has reached $1x10^{20}$ for LTO-9 and enterprise tape surpassing that of the best disk drive reliability by 3-4 orders of magnitude. Media life for modern tape media using BaFe now exceeds 30 years effectively positioning it for archival data. These and several other technological improvements give modern tape significant TCO, carbon footprint and throughput advantages compared to disk. Recent studies project as much as a 75-80% TCO reduction is typical for tape over disk for equivalent storage. Furthermore, the LTO roadmap was recently extended to LTO-14 delivering up to 1,440 TB (1.44 PB) per tape projecting a long life for the LTO family.

A New Formula for Tape Performance Arrives *[High Performance = LTFS + LTO-9 + oRAO]*

The arrival of 18 TB native capacity, LTO-9 drives in Sept. 2021 delivered a higher level of tape performance with the fastest tape data rate and a new feature called oRAO. Open Recommended Access Order (oRAO) was developed by IBM file access acceleration technology and introduced as a native LTO-9 drive function setting it aside from prior LTO generations. oRAO is a data retrieval accelerator that enables tape applications to accelerate the retrieval of files from a single tape by minimizing the seek time (physical tape movement) between files. oRAO supports compressed and uncompressed data, can reduce LTO-9 file access time (time to first byte) up to 73%, and improve resiliency for faster file recovery times (RTO). The formula of a 400 MB/sec. native LTO-9 data rate, combined with oRAO and LTFS access time enhancements are compelling and positions LTO-9 as the fastest performing tape drive available.

***Key point:* Tape has much improved access times, is cheaper to operate, is more reliable, has faster data rates and a much longer media life than any of its disk counterparts.**

Active Archives Benefit From LTFS

The difference between an archive and an active archive is performance. An active archive is a high-value, scalable storage architecture providing faster access to archival data across a virtual file system and manages data across one or more storage technologies. The Active Archive ecosystem can integrate SSDs, NAS nodes, disk, tape, or the cloud (public, hybrid or private). File, block, object, and unstructured formats can be managed with intelligent data movement software and LTFS to provide a highly functional and cost-effective tiered archive solution. New AI and ML applications are reawakening archival data pushing the Active Archive to becoming a de-facto standard storage tier in the future. An Active Archive coupled with the LTO-9 performance enhancements and LTFS and can create an unprecedented archive solution. 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 vendors end user organizations on the new technologies that enable the most effective access to their archived data.

Key point: ***Active archive tiered solutions combining LTFS, disk, tape and SSDs are gaining momentum providing faster access, faster search, and more cost-effective retrieval of long-term archival data.***

LTFS Functionality Summary	
Operating Systems	Cross platform for Linux, Microsoft Windows, Apple OS X (Mac).
Tape Mgmt. SW	No additional tape management, backup software or utilities are required.
Functionality	Functions such as File Open, Drag, Drop, Copy, Paste, Write, Read, Append, Delete and Close from an application are supported on LTFS enabled tape.
LTFS Software Cost	Contact the specific LTFS vendor for any pricing details (typically free).
Where to get LTFS	Download LTFS from one of three vendor's websites (HPE, IBM, Quantum) from the following vendor URLs: • For HPE StoreOpen and Linear Tape File System (LTFS) Software • For IBM Spectrum Archive • 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 of LTO-1 through LTO-4 drives do not support LTFS.
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 applications. **Source: Horison, Inc.**

Key points: ***LTFS implementations can be as simple as a single LTO or enterprise tape drive, small, large or exascale libraries, or part of an integrated tiered storage system with SSD and disk.***

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 certain application deployment. The arrival of LTFS set a new standard for ease of use and portability for LTO and enterprise tape and has opened the door for many new applications. As future tape drive improvements and cartridge capacities continue to arrive, the requirement for faster access times and data retrieval capabilities will become increasingly critical. The stage is set for tape to combine hardware and intelligent software leveraging the LTFS ecosystem to continue to innovate, develop and deliver the most cost-effective secondary storage and archival solutions demanded from the zettabyte era.