

LTFS Takes Tape to the Next Level



Fred Moore, President
Horison, Inc.
www.horison.com

With this long overdue tape access capability now available, LTFS and tape partitioning enable many new uses for tape that weren't practical before.

Introduction

Since its inception, tape has been a sequential access storage device. Though this proved beneficial for “data streaming applications”, the ability to randomly navigate a tape cartridge to locate a specific file or object was physically impossible. With the announcement of LTFS (Linear Tape File System), the long standing rules of tape access have changed. The LTFS Single Drive Edition format provides a new partitioning capability for tape that was first announced on April 12, 2010, and presented at the National Association of Broadcasters (NAB) annual trade show. The LTFS format specification and LTFS file system implementation is supported by IBM, HP, Quantum, and by the LTO Consortium roadmap which defines the future specifications for the LTO tape drive family. LTFS is built on the LTO-5 tape format standard for cartridge interchange with LTO-5 certified tape libraries and drives. In addition to initially providing LTFS tape partitioning for LTO-5 drives, LTFS support is now available for the latest enterprise class tape drives from IBM and Oracle further validating its value. With LTFS, the traditional longer, sequential search times for tape have given way to more disk-like access using familiar drag and drop techniques.

How Does LTFS Work?

The LTO-5 LTFS specification is open source and provides the ability to partition a tape into two segments. With LTFS, the smaller first partition is used to store index information and metadata about the contents of the other, larger user data partition. Metadata is the descriptive information describing user data stored on a cartridge, and it enables searching and accessing of files in the second partition via a graphical user interface provided by the operating system. LTFS creates a self-describing partitioned tape for the first time, meaning that a specific application is no longer needed to figure out what's on the tape because the contents of the second partition are described by the index and metadata in the first partition. The server running the LTFS software graphically displays each cartridge as a subfolder of the library's contents using a GUI format. For a tape library, it keeps a virtual index of the data on all cartridges that are presently in or have been in the library allowing users to search without physically moving or mounting tapes thus improving access time. When a new tape cartridge is mounted, or the data on an existing tape is updated or changed, the LTFS server index is updated with that cartridge's information first, then that index is backed up and written to the cartridge.

LTFS represents a long-awaited breakthrough in tape access providing:

- A file system implemented on dual-partition LTO-5 tape drives and LTO-5 libraries
- Makes tape file access on LTO-5 tape similar to a hard disk, CD/DVD-R disc or USB memory
 - Files and directories show up on desktop as a directory tree listing
 - Enables tagging of files with any text, allowing for more intuitive searches of cartridge and library content
 - Uses the familiar drag-and-drop to move files to and from LTFS tape
 - Can now run many applications on tape that were written to use disk files
 - Supports data interchange between users and systems
 - File Open, Write, Read, Append, Delete and Close functions from an application are supported
- Operating system support (Drivers are available for Apple McIntosh OS X, Windows 7.0, Windows Server 2008 R2, and Linux). *Note: The Mainframe operating system (System Z) does not currently support LTFS*
- Enterprise tape drive support (IBM TS1130, IBM TS1140, Oracle T10000C)
- LTFS is available as open source from HP, IBM and Quantum and can be downloaded from the following URLs
 - For HP <http://www.hp.com/go/LTFS>
 - For IBM <http://www-03.ibm.com/systems/storage/tape/ltfs/browse.html>
 - For Quantum <http://www.quantum.com/ServiceandSupport/OpenSource/LTFS/Index.aspx>

LTFS Extends to Tape Libraries and Enterprise Class Drives...

LTFS was originally designed to improve the management of individual tapes. On July 12, 2011 IBM announced LTFS Library Edition V2.1 extending the capability of the IBM LTFS Single Drive Edition to include the TS2900, TS3310, TS3100, TS3200 and the large-scale TS3500 library system using LTO-5 tape drives. LTFS Library Edition enables mounting of the entire tape library as a file system, while the file manager performs the automation and provides operating system-level access to the entire contents in the library. Since LTFS Library Edition is based on the LTFS format specification, it makes library cartridges 100% interchangeable with cartridges written by single drive versions of LTFS. More importantly, it provides for caching of tape indices, and for searching, querying, and displaying all tapes' contents within an entire IBM tape library without the requirement to physically mount tape cartridges.

On Nov.8, 2011 IBM announced LTFS Library Edition V2.1.1 for IBM LTO-5 Tape Libraries with support for the IBM System Storage TS1140 Enterprise Tape Drive in the IBM TS3500 Tape Library. On-disk metadata storage support, Data Path Failover on selected operating systems, and additional LTFS tools were also announced. At this point, IBM is the only vendor providing LTFS support for its tape libraries. LTFS support is now available for enterprise class drives including the IBM TS1130, IBM TS1140, and the Oracle T10000C. These enhancements signal the increasing value of tape file partitioning as LTFS functionality now extends to the enterprise storage market.

Designed to address these needs, LTFS Library Edition offers:

- Metadata management as a key feature which inherently makes it easier to display all the metadata describing the contents of a tape library
- The ability to recognize the cartridges at library mount time and create a file system directory for each cartridge in the library
- The ability to manage scheduling and accessing tape cartridges across multiple drives in a library

- The ability to pre-populate the file system directories for the cartridge with information from the cached index at library mount time, allowing for listing and searching of the contents of all known tape cartridges without having to physically load each cartridge into a tape drive

LTFS Scorecard – What You Should Know

Cartridge Contents Inquiry for Fast Access and Retrieval	After the tape library is initially inventoried, tape cartridges do not have to be remounted to retrieve content index data and basic cartridge information like volume name, serial number, etc. saving time
LTFS Software Cost	The single tape drive version of the LTFS software is open source and there is no charge for it. Please check with the specific tape library vendor for LTFS Library Edition pricing and details
Where to get LTFS	Download LTFS software from vendor’s website
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 content partition on the tape for backup in case the primary copy is unavailable
Data Sharing Capability	Data stored on tape can be shared between different applications
LTFS Support for Prior and Future LTO Generations	Tape partitioning was introduced with the LTO format in generation 5, and therefore, earlier generations such as LTO-3 and LTO-4 tape drives do not support LTFS. The LTO roadmap indicates that tape partitioning will be supported in all future LTO drives
Operating Systems Supported	Apple McIntosh OS X, Windows 7.0, Windows Server 2008 R2, and Linux
Tape Drives Supported	LTO-5, IBM TS1130, IBM TS1140, Oracle T10000C
Tape Management Software	No additional tape management software or utilities are required

Source: Horison, INC.

Colossal Content - Applications Drive Massive Growth in Digital Data

The era of colossal content (aka big data) is here. New digital data is being generated at over 40% annually and is now being generated by billions of people, not just by data centers as in the past, mandating the emergence of an ever smarter information infrastructure. It is estimated that 75% of all digital data is generated by individuals, not by large businesses; however the majority of this data will eventually wind up in a large business or service provider. In addition, widespread usage for business analytics has significantly increased the value of data stored for long periods of time (digital archives) and much of this data will optimally be stored on tape given the significant TCO benefits and numerous advantages of tape over disk. A comprehensive TCO study by ESG (Enterprise Strategies Group)¹ comparing an LTO-5 tape library system with a SATA disk VTL system for backup using de-duplication shows that disk using deduplication has a 2-4x higher TCO than the tape system for backup in several use cases. Other studies comparing tape and disk for archiving conclude that disk has a TCO up to 15x higher than tape. The TCO advantage for tape is expected to become even more compelling in the foreseeable future. ¹ [A Comparative TCO Study: VTLs and Physical Tape](#), By Mark Peters, ESG, Feb., 2011

With archival and long-term data storage requirements growing at over 60% annually, many businesses are wrestling with the decision on whether to build their archive strategy on disk or tape, or a combination of both. In any case, this growth in digital data is becoming the key demand driver for tier 3 storage and especially for tape. According to IDC’s latest Digital Universe Study of June, 2011, the

world's digital data is more than doubling every two years with approximately 1.8 zettabytes to be created and copied or replicated in 2011 and projected up to 7.9 zettabytes (1×10^{21}) in 2015. According to the study, the magnitude of 1.8 zettabytes of data is equivalent to every person in the world having over 215 million high-resolution MRI scans per day. The table below highlights many additional applications that can potentially benefit from tape, LTFS and its much improved file access capability.

Colossal Content (Big Data) Applications - Key Demand Drivers for Tape and Tier 3 Storage

Digital Assets
Traditional backup and recovery
E-mail archives, ~80% of total e-mail capacity is stored in attachments
Archival Data - Compliance & Litigation with long-term storage requirements
Medical Patient Data, Files/Fixed Images
Insurance Claims and Images, Stock Market Transactions, Banking Records
Web Content, Static Images, Digital Photos
Social Networking Explosion, Wikis
Documents, Printed Materials, Books, Contracts
Rich Media (3D, Multi-dimensional)
Security/ Digital Surveillance/ Forensics
Digital Audio & Video (iPods, MP3, YouTube), Sports Video
Games, Music, Movies, Shopping, Super- imposed Motion Imaging
Medical Images (3D MRIs, CAT Scans, Facial Recognition, Forensics)
Geophysical, Geospatial, GIS, Seismic, Google Earth etc.
Fixed Content and Archival Future Opportunities...
Automated Tiered Storage for Tape and Disk (HSM-like policy-based software)
Tape Partitioning Enhancements for Continually Faster Tape File Access
Cloud Archiving Services Using Tape for Lower Cost, Highly Secure Archiving Solutions

Source: Horison, INC.

For many applications, the lifetime for data preservation and retention has become infinite (meaning the data won't ever be deleted), further increasing the overall size requirements of the digital archive. The growing list of government, compliance and legal regulations worldwide describe the way data should be managed, protected and how long it should be stored throughout its lifetime further swelling archive requirements. For example, the retention period for certain medical records such as X-rays and MRI images need to be kept for the lifetime of the patient though they may only be infrequently, if ever, referenced after they are generated. LTFS has initially generated considerable interest in the media and entertainment industry for storing, editing and accessing video, movie, broadcast and audio files however that is only the beginning. ISVs and end-users can also use APIs (Application Programming Interface) for LTFS to improve and customize workflows for a variety of applications. Customers and vendors are implementing LTFS in a variety of ways, see the following examples.

Customer LTFS Example - Video Using LTFS for Cloud Archiving

Thought Equity Motion based in Denver, Colorado is expanding its hosted video storage to offer an enterprise scale video supply chain as a cloud service. The firm deployed the IBM System Storage TS3200 Tape Library with LTFS as a low cost interoperable, flexible and efficient archive solution. The project involves scanning more than 150 million discrete DPX (Digital Picture Exchange) files and storing them on LTO-5 using the LTFS format. This video is a good example of providing tape archiving services

in the cloud, something that will become more widespread in the future. Their application is described in the video link below. <http://www.youtube.com/watch?v=M7w0jrkQnj4>

Vendor LTFS Example – Crossroad Systems Delivers a New Model for Low-cost File Archiving

Crossroads Systems, Inc. of Austin, TX has introduced the StrongBox architecture that leverages the LTFS technology aiming to take tape-based storage to the next level for both end-users and cloud service providers. Designed to address the need for increasingly lower cost and highly reliable long-term data preservation, StrongBox presents itself as a traditional network-attached storage (NAS) while leveraging the LTFS tape and LTFS architecture for the file storage. Traditionally, accessing files stored on tape often required first reading the whole record to disk and then transferring the data to the host. With LTFS and Crossroads StrongBox, files can be accessed directly from the tape media. No application modifications are required and a persistent view for all files on tapes stored in the library is finally possible. StrongBox combines disk performance with the cost efficiency of tape. <http://www.crossroads.com/>

Future Directions for Tape Storage

Tape innovation continues to progress. Beginning with LTO Ultrium generation 3, WORM (Write-Once-Read-Many) tape was introduced helping to address security and compliance requirements. Tape drive encryption was added starting with LTO generation 4 to provide additional security to protect sensitive data. Beginning with LTO generation 5, LTFS becomes the gateway to a new world of possibilities enabling users to deploy highly cost effective tape solutions for a variety of applications.

Key Factors That Define the Future for Tape

<i>Reliability Now Better for Tape Than Disk</i>	T10000x	1 x 10E ¹⁹ bits	Tape drive BER
	TS11xx	1 x 10E ¹⁹ bits	
	LTO-5	1 x 10E ¹⁷ bits	
BER (Bit Error Rate) Hard Read Errors per Bits Read	Enterprise FC/SAS	1 x 10E ¹⁶ bits	Disk drive BER
	Enterprise SATA	1 x 10E ¹⁵ bits	
	Desktop SATA	1 x 10E ¹⁴ bits	
<i>Capacity and Data Rate</i>	The increasingly popular LTO tape drive family is now in its 5th generation with LTO-5 and 3 future generations are now defined in the LTO roadmap. LTO-5 media has a native capacity of 1.5 terabytes (3.0 terabytes compressed 2:1) and a data rate of 140 MB/sec (280 MB/sec compressed). Today, over 85% of all tape drives shipped are LTO.		
<i>Capacity growth rates</i>	Roadmaps favor tape over disk with 35 TB capability now demonstrated by Fujifilm and IBM.		
<i>Security</i>	LTO-5 tape drives offer both WORM and encryption for higher levels of security and data protection.		
<i>Long Life Media</i>	The shelf life for all new enterprise tape media and LTO-5 media is rated at 30 years or more. Disk typically lasts 3-5 years before replacement.		
<i>Minimize tape file search performance</i>	LTFS partitions the tape into two distinct, individually addressable, unequal partitions with the first quick read partition containing descriptive metadata that enables the quick search capability (random-like) of the data contained in the second partition.		
<i>Inactive data does not consume energy (green)</i>	This is becoming a goal for most data centers for archival data - "If data isn't being used, it shouldn't consume energy" favoring tape over disk.		
<i>Tape TCO Better Than Disk</i>	TCO favors tape for backup (4:1) and archive (15:1) over disk.		

Source: Horison, INC.

The improved management and faster tape file access capability provided by LTF5 becomes increasingly important as tape cartridge capacities and the number of files per tape steadily increases. Fujifilm and IBM have already demonstrated the capability to develop a 35 TB cartridge using barium ferrite media signaling a long and bright future for advanced tape technologies. In addition, several new tape technologies were developed as a result of this joint effort including improved precision control of read-write head positioning, more than a 25-fold increase in the number of tracks, new detection methods to improve read accuracy, and a new low friction read-write head. These enhancements represent significant steps toward achieving areal densities for tape of 100 billion bits per square inch and beyond.

Conclusion

In recent years, tape has improved its position as the preferred choice for the tier 3 data storage explosion. With the advent of LTF5, tape behaves more like disk and is more cost effective than disk. Tape offers a significantly lower total cost of ownership, better reliability, longer life media, faster streaming performance and significantly lower energy costs than today's disk alternatives. As a result of this progress over the past 10 years, the tape industry is aggressively re-positioning itself to address the many new high growth tier 3 storage opportunities which now represent more than 65% of the world's stored digital data, and that percentage is increasing. The latest technology improvements in the tape industry suggest tape is and will continue to be more cost-effectively suited than disk for the enormous archival and colossal content opportunities that lie ahead, whether on-site or in the cloud. Throughout its history, tape has steadily evolved providing faster data rates and higher capacities but the enhancements were always built on sequential access architecture. With LTF5, the traditional tape evolutionary model and has been disrupted changing the rules for tape access while opening the door for a new, more cost effective storage solution than ever before.

End of report

About the author: In 1998, **Fred Moore** founded Horison Information Strategies, a data storage analyst and consulting firm in Boulder, CO. Prior to Horison, Fred began a 21 year career at StorageTek as the first systems engineer and concluded as corporate vice president of Strategy and Marketing. He is a sought-after speaker at IT industry events worldwide and has published numerous reports and publications for the data storage industry.

Email: fmoore@horison.com