

Linear Tape File System by Tandberg Data

Randy Kerns
April, 2011

Technology Insight Series

Evaluator Group

enabling you to make the best storage decisions



Evaluator Group

Introduction

The Linear Tape File System (LTFS) from Tandberg Data provides a file system interface for storing, retrieving, and managing files on LTO-5 and later tape cartridges. Having a file system format similar to disk file systems on tape provides a new level of usage that did not previously exist for tape. Familiarity with using files on disk is ubiquitous so the transition to managing files on a tape device would seem like a natural extension. The LTFS software offered by Tandberg Data brings the characteristics of large capacity and very fast sequential transfer with LTO tape to environments that can capitalize on the need to store and retrieve large files with the flexibility of removable media and interchange capability.

LTFS stores a file system index on each LTO tape cartridge. Each tape can be viewed as a self-contained storage element with the file system index and data immediately accessible when a tape cartridge is loaded into a tape drive attached to a server or workstation running Linux or MacOS. Portability between systems is enabled through a standard format with the file system index maintained on each tape.

LTFS Overview

The Linear Tape File System is software installable on Linux or MacOS operating systems (with Windows availability in 3Q2011) as a file system with a set of unique functions. The Tandberg Data LTO-5 tape drive is attached to a server or workstation and an LTO-5 tape cartridge is inserted. If the tape is a new cartridge, it can be formatted for use with LTFS. The details for usage are in the Tandberg Data LTFS User's Guide.

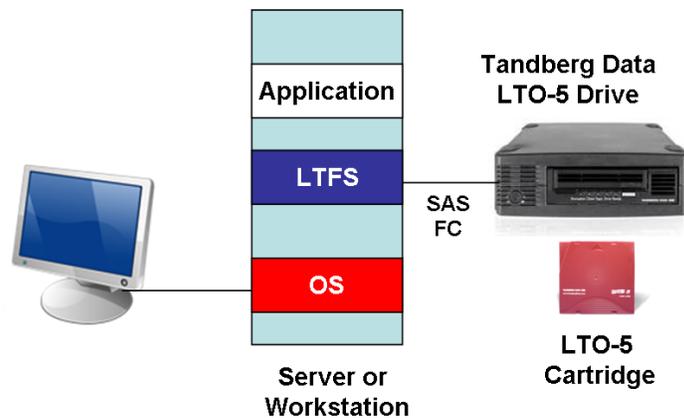


Figure 1: LTFS System Usage

The formatting will create two partitions on the tape: one to store the file system index and the other for data. The file system partition is organized across the length of the tape and reserves approximately 5% of the capacity for the file system index. When a tape with a file system index is inserted, LTFS will present the file information representing the active files on the tape. The tape partitioning is represented in the diagram below.

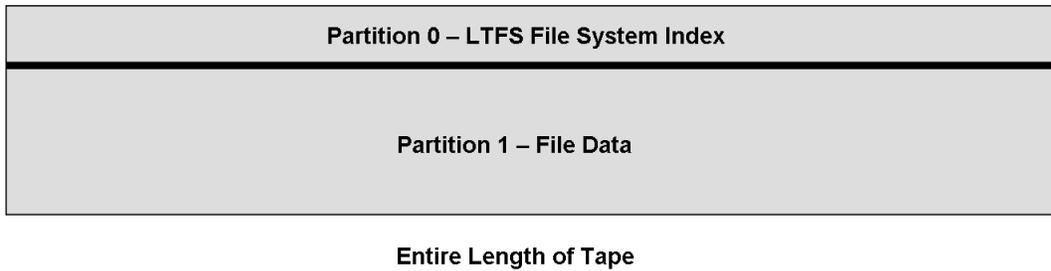


Figure 2: Tape Partitions

Support for two partitions was made available with LTO-5. Tape drives prior to LTO-5 cannot be used with LTFS but future generations of LTO tape will support tape partitions and LTFS. The major highlights with the LTO-5 tape technology are the capacity increase to 1.5 terabytes uncompressed (3.0TB with the usual 2:1 compression expectation) and the performance increase to 140 megabytes/second (280MB/s if the data is compressed). The very high data rate is one of the exploitable benefits with LTFS – large files can be transferred to tape at the high data rate.

Use of LTFS comes in two forms: an interactive Graphical User Interface with the familiar file representation and actions on files such as drag and drop, and a command line interface from which file management functions can be invoked.

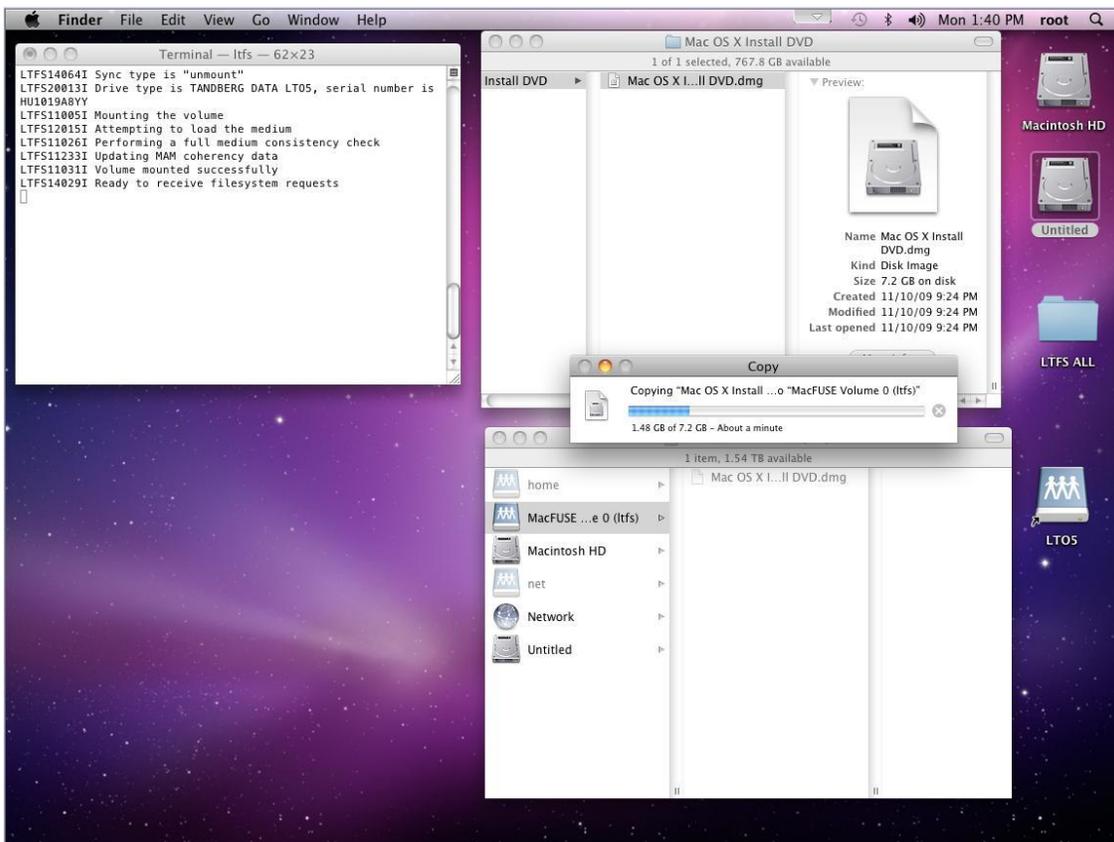


Figure 3: LTFS Graphical User Interface Example

To give a perspective on the file operations using LTFS with a tape cartridge, the following are highlights:

- Formatting a tape into the LTFS format
- Setting file permissions for access control
- Checking the contents of the tape for consistency with the file system index
- Copying a file from a source location (disk) to tape
- Copying a file from the tape to a target location (disk)
- Operational commands such as mount, unmount, and label

Status displays of the tape information and messages based on activities are presented through the GUI as well.

Operationally, when files are written to a tape with LTFS, they are appended to where the end of data was previously. As new files are written, they will be appended until the tape is full (a condition where the file cannot be stored on the tape due to not enough remaining space). A changed file will also be appended with the previous version left in place. The current file GUI will show the newest version but previous versions may be accessed using a command line utility. The following figure illustrates the appending of files and the versioning capability with LTFS.

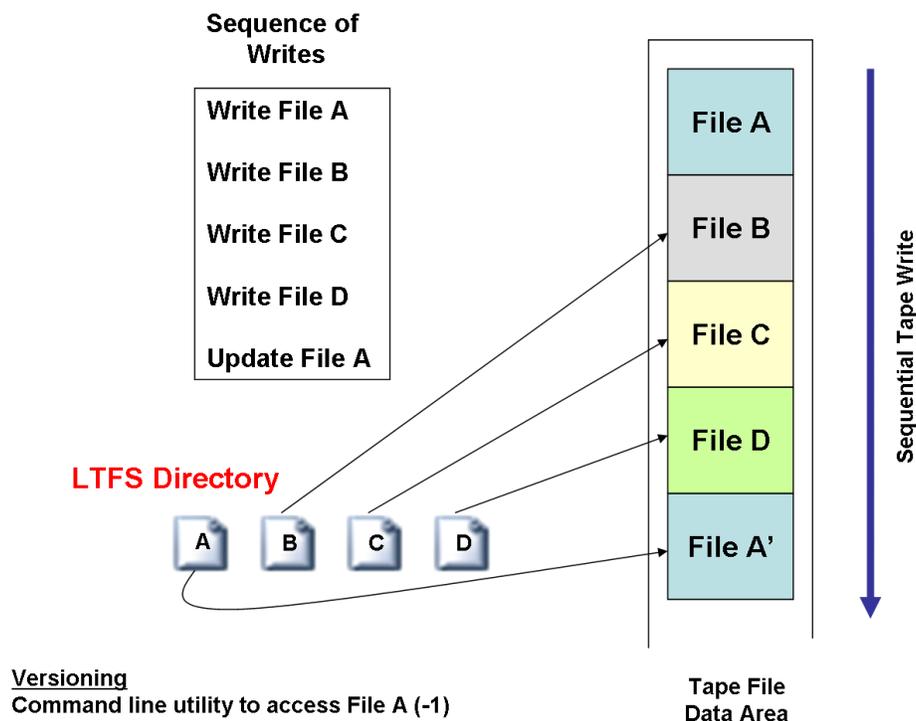


Figure 4: Appending Files and Versioning

Usages with LTFS

The initial focus for LTFS was to exploit the large capacity and high bandwidth of the LTO tape technology. Using LTFS for storing large video files in the media and entertainment market has been an early application. The video files are well suited for LTFS and LTO tape technology: large sequential files, need for interchange with other systems, and high data rate requirements to speed the process. With LTFS, a user can interact directly with the file system by using the GUI to drag and drop files or use the command line to copy files. Applications can also target storing and retrieving data using LTFS. An application would “see” a file in LTFS as a UNC path – a mount point for a file structure or a drive letter.

Beyond the first application usage in the digital video market, other usages characterized by data in large files could use LTFS as the means to store, retrieve, and manage the files on LTO tape technology. Other video usages such as surveillance would have similar requirements. Additionally, capturing log data streams and storing large image files could benefit from LTFS. Applications that collect data for protection and archiving would be able to target tape using LTFS.

Considerations for Usage

One of the very valuable (and cleverly implemented) features of LTFS is the similarity to existing disk-based file systems. Using a Tandberg Data LTO-5 tape drive and tape cartridges with LTFS, a user familiar with disk can immediately store and retrieve files on a mounted tape drive. There are some considerations for a sequential access tape that might be overlooked because of the similarity to disk usage. These should be considered in order to make the most effective use of LTFS:

- The location of the file will depend on when it was written. Because tape is a sequential access media, the time to position to the file and begin transfer will vary.
- The implementation of LTFS requires a single Tandberg Data tape drive attached with a single user. Only one user or application may be accessing the tape at a time.
- Because the index is updated and data is appended on a write or update of a file, the tape cannot be in WORM mode. Additionally, the Secure Erase function is not supported.
- With the management of individual files and the highly-valued interchange capability, tape drive-level encryption is not supported. Files can be encrypted by other means before sending them to the tape through LTFS.
- One operational note to consider is that copying a file from one directory or folder structure located on the tape to another directory or folder structure would cause an inordinate amount of tape motion. A disk intermediary location should be used if there is a reason to do this type of copying of a file.
- File explorers that do pre-reads of files should be avoided for performance reasons. Along the same guidelines, thumbnails should be disabled.

More detailed information is available in the Tandberg Data LTFS User’s Guide.

Review of Benefits

As a quick review, the top-level benefits of using LTFS with Tandberg Data LTO-5 tape technology include:

- LTFS presents a file structure type of interface for managing files on tape. The file structure presented is familiar to users as it is similar to disk-based file systems. The tape contents represented by the file system index are available when a tape is loaded.
- Information interchange between users and different systems is assured with a standard format for storing information and a standard file index structure. The removability of tape cartridges enables mobility and information transfer for large capacity and secured operations.
- The high bandwidth and large capacity of the LTO-5 technology can be used for applications with large file requirements.
- Future generations of LTO tape beyond LTO-5 will operate the same with LTFS.

Future with LTFS

As with many new products seen as highly valuable, product extensions that add more capabilities are quickly identified. With LTFS, a number have been called out as having great potential and are being planned:

- Library support, where tapes in a library could be selectively recalled would expand the usage of LTFS. Establishing a capability to maintain an index across multiple tapes would enable the selection.
- An advanced Graphical User Interface with capabilities such as real-time information about operations in progress (such as moving to the file location on tape) and warnings about un-recommended operations would enhance the experience for the casual or untrained user.
- Additional application software integration to utilize LTFS and manage files stored on tape will provide greater value. Applications that would seem to be immediate, high-value exploiters of LTFS would be:
 - Data protection software
 - Archiving system software
 - Vertical market-specific software where data characteristics such as large files and high transfer rate requirements are needed.

Impact of LTFS

With Tandberg Data making LTFS available at no charge for customers of their LTO-5 tape drives, there are major, positive impacts in using tape technology in open systems. Fundamentally, there is a change in tape usage. The use of tape for archiving of files and for interchange just became much easier and

more useful. Being relegated to storing an encapsulation of data only through a particular application is no longer a hindrance. New usage cases will increase the potential of tape beyond the backup space. LTFS provides a platform for file-based applications to use in established workflows.

The level of user visibility and control of files on tapes opens usage cases and establishes a new relevancy for tape in areas where the large capacity and high bandwidth features are required.

Summary

Tandberg Data provides a new enablement for tape usage with LTFS and LTO-5 tape drives. The immediate value of LTFS has been reviewed here. The extension of usage with new applications and new capabilities will be built on the LTFS foundation available today. With LTFS, more value can be realized from an investment in LTO technology.

###