

Making Everything Easier!™

Oracle Special Edition

LTFS

FOR DUMMIES®

A Wiley Brand

Learn to:

- Manage tape storage like disk storage
- Simplify file access for your users
- Lower your storage TCO

Brought to you by

ORACLE®

Lawrence C. Miller, CISSP



LTFS
FOR
DUMMIES[®]
A Wiley Brand

Oracle Special Edition

by Lawrence C. Miller, CISSP

FOR
DUMMIES[®]
A Wiley Brand

LTFS For Dummies®, Oracle Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2013 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Oracle is a registered trademark of Oracle and/or its affiliates. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-118-81066-8 (pbk); ISBN 978-1-118-81081-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Contents at a Glance

Introduction	1
Chapter 1: Recognizing Today's Digital Storage Challenges	3
Explosive Data Growth.....	3
Rising Costs, Limited Budgets.....	4
Data Movement Challenges	5
Compliance and Accessibility Concerns.....	5
Chapter 2: Understanding What LTFS Is.....	7
Defining LTFS.....	7
LTFS Open Edition and Library Edition	10
LTFS and Oracle's StorageTek T10000 Tape Drives ...	14
Oracle Support for LTFS	14
Chapter 3: How LTFS Addresses Tape Issues and Concerns.....	17
Tape Operations Are Perceived to Be Complex and Cumbersome.....	17
Proprietary Applications Restrict Usability	18
Data Portability and Interchange.....	20
Chapter 4: LTFS Applications and Use Cases.....	23
Simplifying Backup and Recovery	24
Introducing "Tape NAS"	25
Exploring Use Case Examples	26
Taking a Closer Look at M&E Industry Workflows ...	28
Chapter 5: What the Future Holds for LTFS	33
LTFS Challenges	33
LTFS Allows Tape to Be Considered for New Uses ...	35
Chapter 6: Seven Great Resources for Learning More about LTFS.....	37

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz or visit www.wiley.com/go/custompub. For details on licensing the *For Dummies* brand for products or services, contact [Branded Rights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com). Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Vertical Websites

Senior Project Editor:

Zoë Wykes

Editorial Manager: Rev Mengle

Acquisitions Editor:

Connie Santisteban

Senior Business Development

Representative: Karen L. Hattan

Custom Publishing Project

Specialist: Michael Sullivan

Composition Services

Senior Project Coordinator:

Kristie Rees

Layout and Graphics:

Jennifer Creasey

Proofreader: Robert Springer

Special Help from Oracle:

Scott Allen, Michael Brown,
Chris Ilg, Allison Roth,
Carlos Soares, Amy Thompson,
Brian Zents

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher

Andy Cummings, Vice President and Publisher

Mary Bednarek, Executive Director, Acquisitions

Mary C. Corder, Editorial Director

Publishing and Editorial for Consumer Dummies

Kathleen Nebenhaus, Vice President and Executive Publisher

Composition Services

Debbie Stailey, Director of Composition Services

Business Development

Lisa Coleman, Director, New Market and Brand Development

Introduction



For decades, tape storage has remained a viable, cost-effective alternative to disk storage for archiving, backup, and nearline storage requirements. LTFS-OE (Linear Tape File System-Open Edition) software allows a single tape drive to be treated like any other storage device, such as a thumb drive or memory stick. LTFS-LE (Library Edition) software extends LTFS to an entire library, enabling you to manage hundreds, or even thousands, of tape cartridges — petabytes of data — through a basic desktop Explorer interface, as if you were connected to the world's largest thumb drive or NAS device for a fraction of the cost!

LTFS makes tape storage more practical for many new industry uses and applications. *LTFS For Dummies*, Oracle Special Edition, explores IT storage challenges, explains the LTFS specification and different use cases, and gives you a look at how LTFS can help you solve your IT storage challenges today and for the future!

About This Book

Here's a brief look at what awaits you.

Chapter 1: Recognizing Today's Digital Storage Challenges. The book starts with a look at some common storage challenges and how LTFS addresses these challenges.

Chapter 2: Understanding What LTFS Is. This chapter gives you a look at the LTFS specification and how it has evolved.

Chapter 3: How LTFS Addresses Issues and Concerns.

Here, you see how LTFS can help you address your data storage challenges.

Chapter 4: LTFS Applications and Use Cases. In this chapter, you find out about some specific applications and industry uses for tape storage and LTFS.

Chapter 5: What the Future Holds for LTFS. This chapter examines some of the challenges and opportunities ahead for the storage industry.

Chapter 6: Seven Great Resources for Learning More about LTFS. Finally, in that classic *For Dummies* style, you get information about where to go for more about LTFS.

Icons Used in This Book

This book occasionally uses icons to call attention to important information that is particularly worth noting. Here's what to expect.

REMEMBER

This icon points out information that may well be worth committing to your noggin!

TECHNICAL STUFF

This icon explains the jargon beneath the jargon.

TIP

This icon points out helpful suggestions and useful nuggets of information.

WARNING!

These helpful alerts offer practical advice to help you avoid making potentially costly mistakes.

Chapter 1

Recognizing Today's Digital Storage Challenges

.....

In This Chapter

- ▶ Understanding the data growth problem
 - ▶ Working with smaller budgets
 - ▶ Looking at cloud storage and data transportability concerns
 - ▶ Addressing compliance issues and future access
-

In this chapter, you find out about data storage challenges and drivers, and how the Linear Tape File System (LTFS) addresses these challenges in enterprise and tiered storage environments.

Explosive Data Growth

Digital data continues to grow at a stunning rate, and while all enterprises and industries are feeling the growing pains, there are some industries that are experiencing well above average growth. For example, the Media and Entertainment (M&E) industry's storage requirements are exploding with the rapid innovation in video formats, and an hour of professional video

footage from a single camera angle is a struggle to handle on 3TB disk drives. Healthcare, oil and gas, big data, preservation, and video surveillance are all likewise experiencing their own data growth challenges.

Rising Costs, Limited Budgets

Data storage costs have continued to rise, while IT budgets have shrunk for organizations everywhere. Online data storage solutions have traditionally been comprised of disk systems. As the incremental cost of disk storage has risen, many disk vendors have turned to multi-tiered architectures consisting of different combinations of SSD (solid-state drives), SAS (Serial-attached SCSI) drives, and SATA (Serial ATA) drives in an effort to maintain cost effectiveness.

One proven, reliable, and extremely cost-effective storage media that has often been underappreciated is tape. Yet tape has always been a low-cost, high-value storage media, and tape densities and capacities have grown at twice the rate of disk. Innovations in tape technologies have further driven down the cost of tape.



When tape is part of any storage workflow, you can realize significant cost savings. Recent studies by The Clipper Group (www.clipper.com/research/TCG2013009.pdf) have found that disk storage costs 26 times more than comparable tape solutions and the cost gap between disk and tape is growing.

Now, with the introduction of LTFS, tape has become a more user-friendly solution and a viable alternative to disk for nearline storage, in multi-tiered storage architectures for active archiving, and in a growing number of industry-specific use cases.

Data Movement Challenges

The digital storage industry has also seen the emergence of Storage-as-a-Service (SaaS, not to be confused with Software-as-a-Service), or cloud offerings, in recent years. As these solutions gain traction, storage administrators are considering a cloud tier in their storage architecture. One of the challenges with implementing a cloud tier is the movement of data between the cloud and the data center. Fiber networks are an expensive option to move large data sets for disaster recovery or file recalls, and shipping disks is risky due to the possibility of physical or mechanical damage during transport.

However, tape has always been the best technology for transporting large data sets, and with LTFS shipping data on tape has never been easier.

Compliance and Accessibility Concerns

Legal, regulatory, and corporate governance requirements are all driving increasingly longer data retention requirements and creating new data storage and management challenges for IT organizations. These challenges are best solved with tape storage.

Litigation is an ever-present specter for many organizations today. Electronic discovery of digital data can be very costly for these organizations, both in terms of time and legal fees.

Regulatory requirements include the U.S. Health Insurance Portability and Accountability Act (HIPAA),

Sarbanes-Oxley (SOX), and Canada's Management of Information Technology Security (MITS) directive, among others.

Finally, according to the Storage Networking Industry Association (SNIA), 80 percent of organizations participating in a recent survey responded that they are required to retain data for more than 50 years, and 68 percent of companies require a 100-year archive!

Managing vast repositories and archives is a challenge in and of itself. But what computers and technology will exist 10, 50, 100, or even 1,000 years from now to retrieve this data? LTFS is one solution that can help address these challenges today and for the future.

Chapter 2

Understanding What LTFS Is

In This Chapter

- ▶ Learning the basics of LTFS
- ▶ Using LTFS for tape drives and libraries
- ▶ Running LTFS with Oracle's StorageTek T10000 tape drives
- ▶ Getting support for LTFS

In this chapter, you get the lowdown on LTFS and its basic operation. You also learn about the differences between the single tape drive and tape library versions of LTFS. And you take a look at how LTFS performs with Oracle's StorageTek 10000C tape drive.

Defining LTFS

The Linear Tape File System (LTFS) format is an open specification that describes the layout of data structures stored on sequential access media, such as a magnetic tape cartridge.



In computing, an *open specification* refers to a detailed set of nonproprietary, technical requirements for a proposed standard.

LTFS software can provide users with a *file system view* of a tape's content. With LTFS, files can be retrieved directly from tape without a backup/archive application and without staging the data on another storage tier — such as a flash drive or a disk storage system.



A file system view presents a hierarchical, folder-based representation of the contents of a disk or LTFS-formatted tape.

In order to present a complete file image to a user, LTFS stores files in a self-describing format by logically splitting a tape cartridge into two partitions that contain the following data (see Figure 2-1):

- ✔ **File metadata** contains the file structure, file names, file size, file format, and other data elements that are indexed to simplify finding and accessing data. This metadata is what enables any LTFS-formatted tape cartridge to be *self-describing*, meaning it contains data about the data itself. Within the metadata partition, files are stored in a hierarchical directory structure.
- ✔ **File data** is the raw file content that is stored in the same way it has traditionally been stored on tape cartridges.

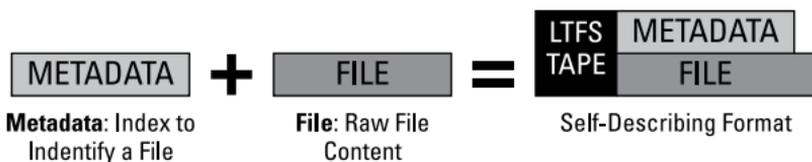


Figure 2-1: A self-describing format contains metadata and file data on the same storage medium.



A self-describing format is a specification for writing descriptive metadata and file data to the same storage medium.

When tape media is loaded into a tape drive, the complete file folder image is displayed. The file structure metadata is pulled from the first partition, and the raw file content is accessed from the second partition (see Figure 2-2).



Figure 2-2: LTF5-formatted tape partitions.



With tape storage, file metadata has historically resided solely within a proprietary backup or archive application making it difficult to access files on tape without the presence of the application (see Figure 2-3).

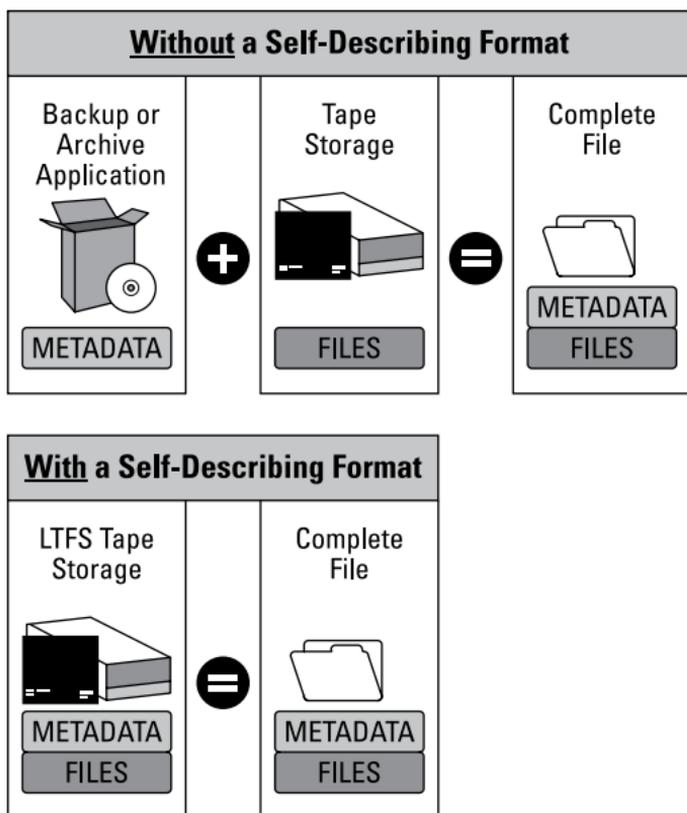


Figure 2-3: Tape storage without a self-describing format versus with a self-describing format.

LTFS Open Edition and Library Edition

The LTFS-OE specification was released in 2010. LTFS enables users to read files from, and write files to, a

tape cartridge with a single tape drive much the same way files are accessed and written to disk and flash storage devices.



In August 2012, the Storage Networking Industry Association (SNIA) formed an LTFS Technical Work Group (TWG) and named Oracle and IBM as the co-chairs of the LTFS TWG. In October 2012, the LTFS TWG published version 2.1.0 of the LTFS specification for public comment and is moving toward adoption of LTFS as an industry standard recognized by the American National Standards Institute (ANSI) and International Organization for Standardization (ISO).

The open format of LTFS addresses transportability issues by allowing organizations to freely exchange tape cartridges and be able to view and manage the contents without the need for proprietary tape software applications.

The single-drive version of LTFS, LTFS-OE, has solved several key issues related to the interchange and portability of data on tape. LTFS-OE has been adopted within several industries where data interchange and portability between field-based sites (such as on-location movie sets or oil and gas exploration sites) and a centralized data processing location are required.

The single-drive limitation of LTFS-OE limits the ability of LTFS-OE to address a broader market. LTFS-OE has no tape volume management, requiring a user to mount a cartridge to see what's on it. Additionally, LTFS-OE has limited scalability because its use is subject to a single tape drive.

To address enterprise storage requirements in multi-tiered storage architectures, active archives, and data protection solutions, an LTFS-based solution for much larger storage environments is needed.

LTFS-LE expands LTFS from a single tape drive to an entire tape library consisting of multiple tape drives and hundreds, if not thousands, of tape cartridges. Using a basic explorer or browser interface, a user will see a tape library as a very large pool of storage — much like a NAS disk storage system — which can be organized according to the users' preferences. Each tape cartridge appears as a folder under the library.

After a tape cartridge is first mounted, LTFS-LE reads and makes a copy of the file metadata and saves it. With a copy of all the metadata for all of the files in the library, LTFS-LE presents the user with a folder containing all of the files on each of the cartridges. LTFS-LE knows which files are on which tapes in the entire library, so that the user doesn't have to manage the files, cartridges, or the cartridge locations in the library (see Figure 2-4).

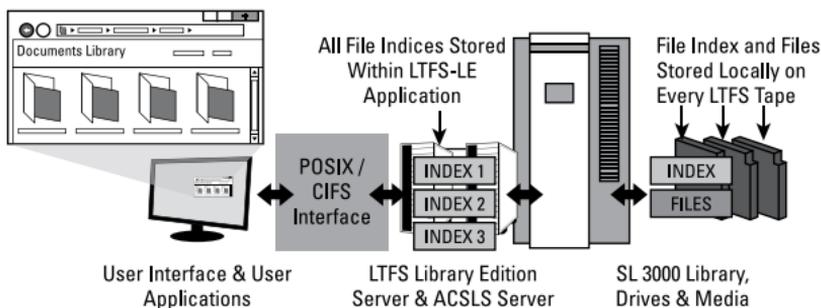


Figure 2-4: LTFS-LE provides visibility into all the files in a tape library.

Some of the features and benefits of LTFS-LE include

- ✔ Simplified tape management where users manage folders and the LTFS-LE software manages the library robotics, slots/cells, cartridges, and tape drives.
- ✔ Files can be transferred seamlessly between media types without the need for proprietary software.
- ✔ Support for enterprise tape drives and 5TB and larger LTFS-formatted tape cartridges.
- ✔ Flexible software supported on multiple server platforms.
- ✔ Easy deployment leveraging a partition of an existing tape library.
- ✔ An open format that supports LTFS Format 2.0, HP, IBM, and StorageTek Drives.
- ✔ Scalable support from entry-level to enterprise modular tape libraries.
- ✔ Easy migration of data to new technology when today's drives reach end-of-life (EOL).
- ✔ Communication with tape storage through standard file system interfaces. No need to modify enterprise applications for tape storage.
- ✔ Ability to create transport copies directly from an archive and easily share copies on tape cartridges.

LTFS and Oracle's StorageTek T10000 Tape Drives

Oracle extends the LTFS specification to enterprise tape drive technology by adding support for the T10000 tape drives. Oracle's StorageTek LTFS software supports a wide range of tape technologies including HP and IBM LTO tape drives, as well as Oracle's enterprise-class StorageTek T10000 tape drives.



With LTFS, files originally written in an LTO LTFS format are easily transferable to Oracle's StorageTek T10000 tape drives.

The StorageTek T10000D drive can store up to 8.5TB of native capacity — the equivalent of 369 hours of high-definition video — onto a single T10000D cartridge.

Oracle Support for LTFS

Oracle Services helps you migrate, deploy, and maintain Oracle's StorageTek tape solutions. These services include

- ✓ **Oracle Premier Support for Software.** Provides access to product updates and enhancements, as well as technical assistance and support resources to help you maintain your product, optimize performance, and effectively implement new LTFS software functionality. Download the data sheet at www.oracle.com/us/support/premier/software/overview/index.html.
- ✓ **Oracle Support Advantage.** Delivers higher value and a better ownership experience for Oracle

customers. Visit www.oracle.com/us/support/premier/engineered-systems-solutions/overview/index.html.

- ✔ **Oracle Advanced Customer Support for Server and Storage Systems.** Provides know-how for deploying and optimizing Oracle hardware systems to improve performance, increase availability, and reduce implementation times. Download the data sheet at www.oracle.com/us/support/advanced-customer-services/servers-storage/overview/index.htm.

Chapter 3

How LTFS Addresses Tape Issues and Concerns

In This Chapter

- ▶ Recognizing interoperability issues
 - ▶ Tackling usability with “open” solutions
 - ▶ Moving data between organizations
-

This chapter takes a look at some of the perceived concerns — some more real than others — associated with tape storage and how LTFS helps to address those issues and opens new possibilities for using tape storage in IT operations.

Tape Operations Are Perceived to Be Complex and Cumbersome

Most IT organizations have a data center environment comprised of numerous systems and applications with various hardware and software components from multiple vendors. Managing the interoperability of systems

and applications through the full hardware and software stack is often a very time-consuming and resource-intensive activity for IT staff.

In particular, IT must dedicate resources to align the processes and technical resources for data exchange mechanisms. These processes also require maintenance and verification, because both application and hardware components change over time. The resulting costs can significantly drive up an organization's capital and operating expenses.

Tape storage is not immune to these complexities. Tape storage is often viewed by IT as just another tier of storage with its own file management software. This makes tape storage appear untenable for many applications.

Proprietary Applications Restrict Usability

Perceived usability concerns exacerbate interoperability issues, which have limited the use of tape as a storage alternative in many application environments.

Conversely, IT organizations have grown to embrace almost anything that can be labeled as "open." To a large extent, open source software has not penetrated the market for tape storage technologies and backup/archive software as prevalently as in other markets and technologies. Instead, IT organizations are at the mercy of their backup or archive vendors' proprietary software because the only way to access backup or archived data

is through these applications. This dependence can be problematic for IT organizations, for example, when a solution becomes end-of-life (EOL) or a vendor simply disappears — perhaps due to bankruptcy or acquisition by a competitor.

When a solution is no longer supported, IT has no easy or inexpensive way to access legacy backup or archive data. Expensive and time-consuming services are typically needed to transfer this data to a newer platform. The need for an open source, open standard, self-describing tape format is clear.

Ultimately, tape needs to perform more like disk by providing a single interface or access point to the operating system and application environment. Prior to self-describing technologies like LTFS, you could not access files directly from tape without expensive and often proprietary backup, archive, or HSM (hierarchical storage management) software.

LTFS simplifies access to files on tape. With LTFS, you can retrieve files directly from tape without the need for a backup, archive, or HSM application. Because LTFS writes data in an open and self-describing format, anyone with a compatible tape drive can download an LTFS driver and read an LTFS tape without an archive application or any other software.

You can access files that are stored in an LTFS format on tape just as if the files were stored on a NAS storage device (see Figure 3-1). This is why LTFS is often referred to as “Tape NAS.”

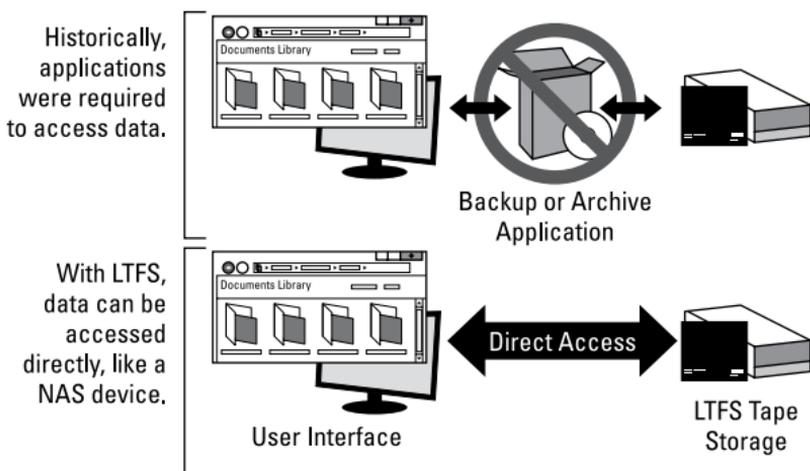


Figure 3-1: LTFS enables easy tape file access without a backup or archive application.

Data Portability and Interchange

Often, the need arises for large volumes of data — from a few gigabytes to hundreds of terabytes or more — to be shared between organizations. Data portability and interchange is necessary when different organizations must collaborate on a project or otherwise have access to the same data, for example, in the Media and Entertainment (M&E) (discussed in Chapter 4) and oil and gas industries, and in video surveillance, scientific research, litigation, disaster recovery, or any other event or activity that requires large data sets to be physically transported.

However, hard drives are extremely fragile and subject to failure during transport, often due to excessive vibrations, extreme environmental conditions, or poor handling. And portable flash storage devices, while compact, are limited in capacity and cost prohibitive for large data transfers.

Tape cartridges are a superior transport option because they are highly durable, very high capacity, and not nearly as expensive as hard disk drives or flash storage (see Figure 3-2).

	Fibre	Flash	Disk	Tape
				
Cost:	\$\$\$\$	\$\$\$	\$\$	\$
Availability:	Limited	Global	Global	Global
Reliability:	High	High	Low	High

Figure 3-2: Tape offers the lowest cost and most reliable way to transfer files.

With LTF5, users can create/capture data at one location and easily and inexpensively send the data to another location for analysis, processing, or long-term storage. Storing files in the LTF5 self-describing format enables users to share content with ease. Thus with LTF5, tape has become the best option for file transport in terms of cost, reliability, durability, capacity, and ease of use.

Chapter 4

LTFS Applications and Use Cases

.....

In This Chapter

- ▶ Recognizing backup and recovery challenges
 - ▶ Switching from disk- to tape-based NAS
 - ▶ Considering different industry uses for LTFS
 - ▶ Using LTFS in media and entertainment workflows
-

LTFS is ideal for many different use cases, including workloads in which tape is a nearline storage platform and organizations that have active archive requirements or are required to exchange large amounts of data in an easy-to-access format. Additionally, organizations that use tape storage for backup, archiving, or disaster recovery, regardless of industry, should consider LTFS technology.

In this chapter, you learn about LTFS applications and some common LTFS use cases.

Simplifying Backup and Recovery

Organizations must retain backup copies or archives of primary data for many years, if not forever, due to various business and regulatory requirements. Many industries are beginning to save more data sets for customer analysis and big data initiatives they may perform in the future. Gleaning unforeseen information and trends from historical and static data is becoming more common across many industries.

Most backup and recovery systems today typically rely on vendor-specific proprietary software for all aspects of backup and recovery including scheduling, encryption, cataloging or indexing, verification, storage formatting, and various other operations. Backup software, like any other software, must periodically be updated to the latest version or replaced altogether.



Proprietary backup or archive software can lock an organization into a specific vendor's solutions if they do not support self-describing formats.

While you can expect some level of backward compatibility among different versions of an individual vendor's software, you are practically guaranteed no compatibility between software from different vendors. Thus, in order to change your backup software, you may need to either keep a copy of your legacy backup software (and any associated backup hardware) readily available indefinitely, or undertake a lengthy and expensive data migration to a new platform. Rinse and repeat every couple of years as technology and your unique business requirements change!

Finally, endless litigation has become an unfortunate fact of life for many organizations throughout the world today. Subpoenas often require organizations to produce true copies of all their backup data within the scope of an electronic discovery request. Here again, outdated versions or legacy proprietary backup software from different vendors can significantly impede an organization's ability to comply with such a court order and cause legal costs to skyrocket.

By abstracting proprietary backup software from the underlying file system, LTFS addresses all of these backup, recovery, and long-term data retention challenges.

Introducing "Tape NAS"

Network-attached storage (NAS) is commonly used to provide organizations with disk storage for file serving. However, disk-based storage, such as NAS, has its limitations as well.

When comparing cost per terabyte, disk storage is significantly more expensive than tape storage, and when comparing total cost of ownership (TCO), disk is "off the charts" more expensive than tape storage. In addition, disk storage is less reliable and has a significantly shorter life span than tape storage. Consequently, for certain workloads including, but not limited to, backup and compliance archiving, tape is the preferred storage medium. Cutting costs by implementing tape can be challenging as many organizations have architected their data centers to leverage disk exclusively.

Now with LTFS, organizations can access tape media in exactly the same manner as disk storage, using a standard application interface such as Windows Explorer. Tape storage with LTFS is positioned to rival disk-based

NAS storage for certain workloads. This is particularly true for data use profiles consisting of long retention periods with relatively infrequent access, such as compliance and archive data. In addition, smaller organizations outgrowing their storage budgets due to disk costs now have a tape NAS solution with LTFS-LE.



With LTFS-LE, you can move data to and from tape by simply dragging and dropping files from one window to another, as you would with disk in most popular operating systems.

Exploring Use Case Examples

Industries and organizations with a clear need for LTFS storage technology and solutions include

- ✔ **Cloud service providers.** As the popularity of cloud services grows, cloud service providers are looking for more cost-effective ways to transport, store, and for data protection purposes, recover bulk amounts of their customer's data. LTFS provides mechanisms that allow third-party cloud providers to more efficiently utilize low-cost, highly reliable tape storage as part of their infrastructure.
- ✔ **Healthcare.** As mandates for medical imaging, electronic medical records (EMR), and electronic health records (EHR) are implemented, the healthcare industry's nearline storage needs grow exponentially. High-definition medical images such as MRIs, X-rays, mammograms, and CT scans consume enormous amounts of storage and by law must be accessible for the patient's lifetime. With each hospital, clinic, and doctor's office using disparate systems, a nonproprietary, open format is required. Transferring large numbers of high-definition

medical images via encrypted tape is often more cost effective and more secure than using dedicated networks. LTFS-based tape storage is a great solution to many of the storage challenges in this industry.

- ✔ **Media and Entertainment (M&E).** M&E companies often write tremendous amounts of data, not only at the original capture point (high-definition digital video camera) but also as files move through the post- production and editing process. Companies want to keep the original file as well as additional edited files in the archive. All of this data may need to be recalled at any given time. In addition, much of this data is captured off-site and needs to be transportable. For example, the average, full-length motion picture shot with digital cameras consumes up to two petabytes of digital data storage before it is finished, almost all of which needs to be retained long term.
- ✔ **Oil and Gas (O&G).** O&G companies explore remote locations in search of new oil reserves. Part of their search involves seismic imaging and ultrasounds to find out what is underneath the surface. Like other industries, O&G companies need transportable storage and a simple, open format to manage files that may be compiled into enormous data sets.
- ✔ **Digital video surveillance.** Similar to reality television shows, digital video surveillance in the private sector, as well as federal, state, and local governments, records hours of security image feeds. The most economical solution would be to immediately archive these feeds on tape for future recall.

Taking a Closer Look at Media and Entertainment Industry Workflows

The M&E industry provides some particularly insightful use cases that are applicable across many industries and warrant further discussion.

The most basic use case emphasizes the need for transportable storage. M&E companies commonly outsource work during the production of a movie. For example, a third party might be responsible for adding a soundtrack or sound effects. The exchanged files can usually be measured in terabytes (TB). For example, according to a January 2010 Tech Target story, the film *Avatar* required 100TB of data storage just for scene-rendering. Today, it is not uncommon for total storage consumed in the making of a feature-length digital motion picture to reach two petabytes.

File sizes for M&E will continue to increase as the industry transitions from 1080P to higher-definition 2K/4K/8K 3-D formats. Due to the costs associated with disk storage and network bandwidth, M&E companies prefer to exchange files via tape storage.

M&E companies also benefit by having transportable, cost-effective storage on the sets of their productions. The Sony XDCam is a market leader for production cameras and can record at 126GB per hour. A reality television show such as *Big Brother* will use as many as 28 cameras simultaneously plus 60 microphones. This

means that 3.65TB of raw video and audio content can be created per hour. As Table 4-1 shows, the expense of recording all of this data can skyrocket depending on the type of storage media used.

Table 4-1 Storage Media Costs for 3.65TB of Raw Video/Audio Content

<i>Storage Format</i>	<i>\$/Hr</i>
SxS Express 34 card (native Sony XDCam media)	\$200,750
146GB SAS hard disk drive	\$5,564
LTFS-formatted LTO tape	\$133

Using LTFS provides the following benefits:

- ✓ Reduces transportation costs
- ✓ Doesn't require compatible proprietary storage
- ✓ Simplifies workflows
- ✓ Can make transportation copies directly from archive
- ✓ Guarantees archive copy immediately upon import when files return from third parties
- ✓ Reduces storage costs — storing content on XDCam is almost 10 times more expensive than storing on digital tape

Figures 4-1 and 4-2, respectively, show examples of how M&E companies exchange files without LTFS and with LTFS.

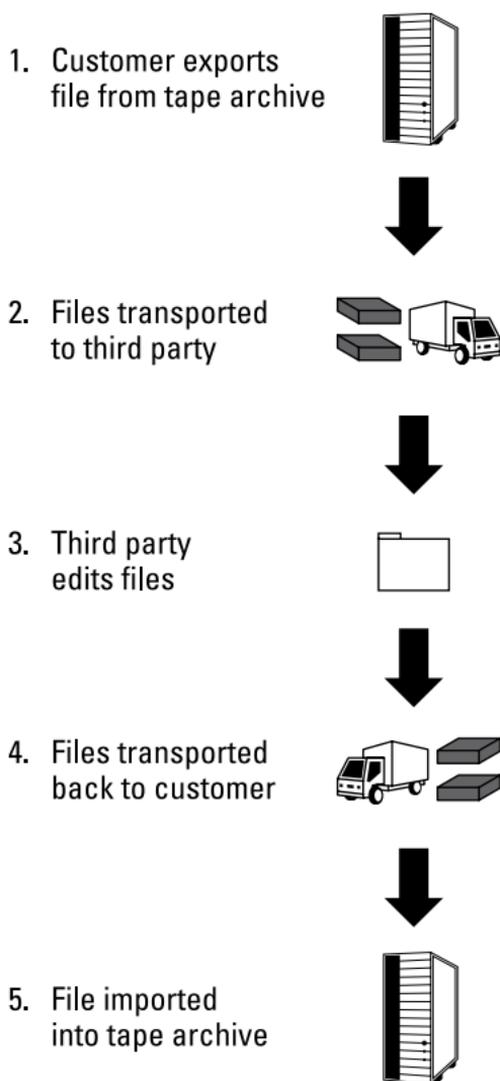


Figure 4-1: M&E workflow without LTFS.

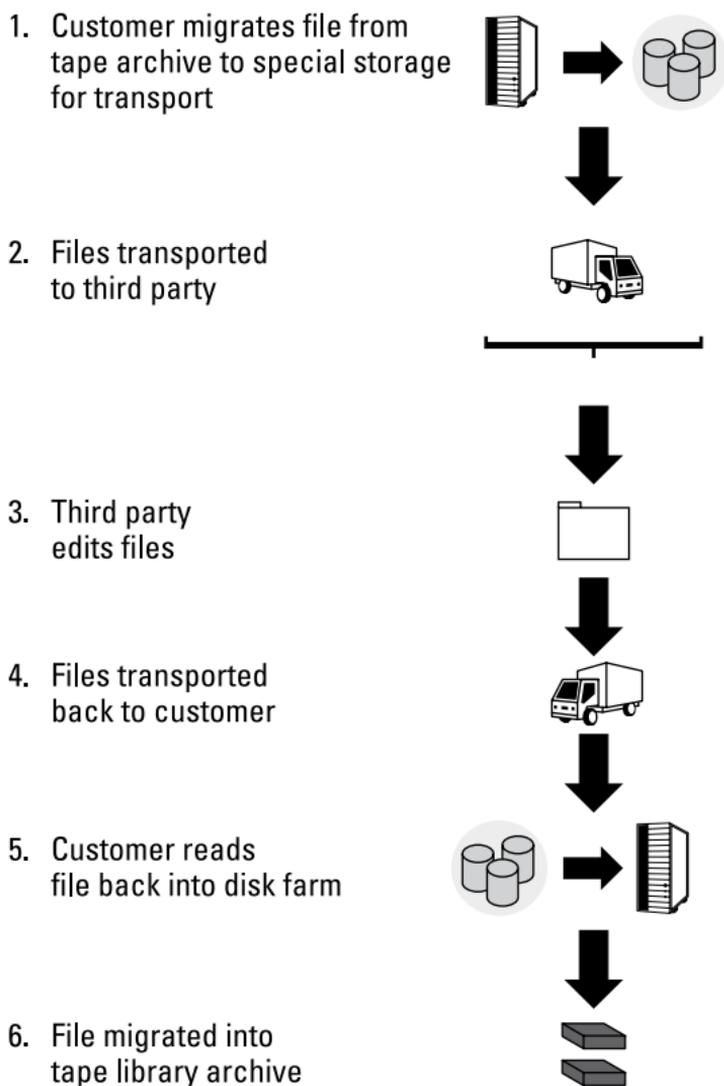


Figure 4-2: M&E workflow with LTFS.

The cost savings from tape storage are only beneficial if the production team has an easy way to access the content. LTFS-formatted tape cartridges are self-describing, so their content can be easily managed for ingest.

In addition, you can create an archive copy of raw footage immediately before sending it to a disk-based editing suite. LTFS then enables an editor to automatically create (by policy) an archived copy of edited content before moving it on to the next step of the production workflow. So ultimately, LTFS can guarantee that multiple archive copies are available, while eliminating extra layers of expensive disk storage.

The richness of the metadata information that M&E companies are able to use with their tape storage will be enhanced as LTFS evolves further.

For example, if a movie file was 15TB, it could be recalled on a pair of Oracle's StorageTek T10000D tape drives (see Chapter 2) at 252MB per second in less than 6.75 hours since drives can be multi-threaded. While this is extremely fast given the volume of data, it may be impractical if only one specific scene is needed rather than an entire movie. However, if the movie is segmented into files on a LTFS formatted tape, then the time required to access the specific segment or scene is dramatically reduced.

In other words, LTFS opens up opportunities to allow existing applications' file-based workflows to leverage the StorageTek T10000 drive's capacity and performance. Similar scenarios are applicable for broadcast and news agencies that need a large archive of nearline storage.

Chapter 5

What the Future Holds for LTFS

.....

In This Chapter

- ▶ Taking a closer look at LTFS adoption
 - ▶ Going (boldly) where no tape has gone before!
-

Many myths and misconceptions about tape storage create challenges for its further adoption in the IT enterprise. Even with capacity shipments continuing to increase each year and the price of tape storage continuing to decline faster than other storage media, storage administrators have sometimes overlooked tape technologies.

This chapter explores some of the tape adoption challenges in more detail and how LTFS brings new life to tried-and-true tape storage!

LTFS Challenges

Perhaps the biggest challenge for LTFS technology is that disk storage vendors view LTFS as a competitive threat.

LTFS neutralizes many advantages of disk storage by allowing traditional tape storage strengths — such as its low cost per terabyte, ultra scalability, greater reliability and durability, easier transportability, and nearly non-existent energy costs — to become even stronger differentiators.

Backup and archive application vendors have also been slow to adopt LTFS thus far, for several reasons, including

- ✔ **Defending existing market share.** LTFS may allow customers to more easily move from one tape software vendor to another or in some cases forgo the proprietary software altogether.
- ✔ **Resisting open source.** For many vendors, “open source” is synonymous with “free” so they have concerns about limited revenue opportunities. Software vendors that treat perpetual licensing of proprietary software as their primary source of revenue are thus slow to adopt new, open source technologies.
- ✔ **Waiting for a standard.** This is perhaps the most legitimate reason for a vendor to hesitate before adopting LTFS. Minor variations between different vendor implementations of LTFS today can lead to later versioning of software and hardware solutions, when LTFS becomes a standard.

Despite these obvious challenges, Oracle and the other members of the LTFS Technical Work Group (TWG), as well as many other storage and industry-specific hardware and software solution providers, are committed to further expanding and enhancing the LTFS format.

Promoting a simplified data interchange between different tape drives — such as LTO tape drives and Oracle’s StorageTek T10000 tape drives — brings organizations another step closer to efficiently managing mixed LTFS media in a single tape library.



It makes sense for organizations to implement LTFS across current tape libraries and LTFS-compatible tape drives, rather than “ripping and replacing” existing storage and archive infrastructures.

LTFS Allows Tape to Be Considered for New Uses

Future LTFS capabilities have some intriguing possibilities.

For example, imagine a world where storage administrators are no longer bound to a single backup or archive application, but rather are able to freely move large data sets from one application to another with ease. As retention requirements change and the perceived value of data sets shift, you need the freedom to select the best application to meet your needs. Software application vendors want you to be bound to their proprietary formats, but as LTFS works its way into the industry, customer requirements are driving smarter applications that provide the flexibility you demand.

In addition, as storage needs continue to increase, retention timelines are also extending from tens to hundreds of years. Building sustainable archives is a real challenge, and data center architects can’t rely on a single archive application. Organizations are demanding

open and self-describing formats to ensure that they can access their data well into the future.

As LTFS nears adoption as an industry standard, more software application developers and systems integrators inevitably will embrace the inherent benefits of this open source, open standard technology and enable new uses and advanced capabilities to address the storage needs of their customers.

Chapter 6

Seven Great Resources for Learning More about LTFS

In This Chapter

- ▶ Getting more information about LTFS
-

Should you want to learn more about LTFS, this chapter points you to several resources that can help you further your knowledge and skills!

Analysis Tools

A tiered-storage architecture consisting of flash, disk, and tape storage working together to optimize a data storage infrastructure, can offer substantial savings over single-tier and dual-tier disk-only approaches (see Figure 6-1). Tiered-storage savings are achieved by using policies to place data on the appropriate storage tier as it moves through its extended life cycle.

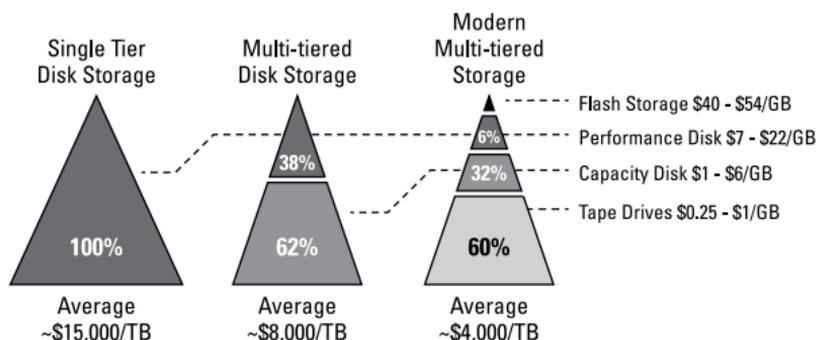


Figure 6-1: Cost comparison across different storage architectures.

Go to www.oracle.com/us/media/calculator/tieredstorage/index.html to analyze your organization's storage architecture and potential savings, using Oracle's free Tiered Storage Savings Calculator.

Documentation

If you're ready to roll up your sleeves and get down to the nitty-gritty details of LTFS, you can download the complete 60-page LTFS Format Specification at <http://docs.oracle.com/cd/E19957-01/LTFSSpec/LTFSSpec.pdf>.

The LTFS Format Specification document includes definitions and acronyms, volume layout, data extents and formats, label and index formats, MAM (Medium Auxiliary Memory) parameters, certification and schema information, and LTFS volume and index examples.

Open Source Software Portal

The importance of using open standards — in both open source and proprietary software — cannot be overstated. Open standards help to ensure interoperability between different vendor products and solutions.

Organizations and enterprises around the world are using Oracle open standard solutions together with other open source technologies in mission-critical environments and are reaping the benefits of lower costs, easier manageability, and higher availability and reliability, along with performance and scalability advantages.

Oracle's Open Source Software (OSS) portal allows you to explore all free and/or open source software from Oracle. Go to <https://oss.oracle.com>.

Oracle Product Pages

You can read analyst reports, customer case studies, data sheets, and press releases about all of Oracle's tape storage products and solutions on Oracle's Product Pages at www.oracle.com/us/products/servers-storage/storage/tape-storage/overview/index.html.

Oracle's StorageTek tape and library solutions include

- ✔ **Tape libraries.** Oracle's StorageTek tape libraries are highly scalable and ensure data availability in heterogeneous environments of any size, including Oracle applications, Windows desktops,

mainframes, and supercomputers. Oracle tape library solutions include

- StorageTek SL8500 and SL3000 Modular Library Systems
- StorageTek SL150 Modular Tape Library

✔ **Tape drives.** Oracle's StorageTek speed- and capacity-optimized tape drives enable robust data protection, energy efficiency, and fast data transfer. Oracle tape drive solutions include

- StorageTek T10000C and T1000D tape drives
- StorageTek LTO tape drives

✔ **Tape media.** Oracle StorageTek tape cartridges ensure media durability, performance, and longevity in even the most demanding environments. Oracle tape media products include StorageTek T10000C, T10000D, and LTO data cartridges.

✔ **Tape software.** Oracle's StorageTek tape management and tape device software and encryption offerings ensure simplicity and security across your tape infrastructure. Oracle tape software and encryption solutions include

- StorageTek Linear Tape File System, Open Edition and Library Edition
- StorageTek Tape Analytics Software
- StorageTek Automated Cartridge System Library Software (ACSL)
- StorageTek Storage Archive Manager

Oracle Technology Network

The Oracle Technology Network (OTN) is the world's largest community of application developers, database administrators, system administrators/developers, and architects using industry-standard technologies, such as LTFS, in combination with Oracle products and solutions.

OTN membership is free and provides online access to millions of members who collaborate to share real-world expertise and best practices about how to design, build, deploy, manage, and optimize applications. OTN is also a provider of free developer workshops around the world for you to learn how to program in Java, and a sponsor of third-party conferences, user group meetings, and other events. Join OTN at www.oracle.com/technetwork/index.html.

OTN Garage

The OTN Garage is the official blog of the System and Developer Community of OTN (see the preceding section). You can read about the latest LTFS and tape storage news and trends at <https://blogs.oracle.com/OTNGarage/>.

In addition, you can read blogs in more than a dozen other OTN Garage communities. You can also follow your favorite OTN bloggers on their personal websites and newsletters, Facebook, and Twitter, and subscribe to RSS (Really Simple Syndication) and Atom newsfeeds.

SNIA Website

The Storage Networking Industry Association (SNIA) is a nonprofit organization consisting of more than 400 member companies (including Oracle), channel partners, and end-users. The goal of SNIA is “to promote acceptance, deployment, and confidence in storage-related architectures, systems, services, and technologies, across IT and business communities.”

SNIA resources include best practices, events, newsletters, social and multimedia, technical communities and working groups, training and certification programs, tutorials, white papers, and much more. Go to www.snia.org/.

What are the analysts saying about disk versus tape?

The gap between the cost of disk and tape solutions is continuing to increase. Learn more about what the storage industry is saying about the cost savings of tape.

Get a free industry report from the Clipper Group at www.oracle.com/webapps/dialogue/ns/dlgwelcome.jsp?p_ext=Y&p_dlg_id=9983736&src=7011670&Act=123.

Access, archive, and transport large files on inexpensive tape storage

The Linear Tape File System (LTFS) is a rapidly evolving storage industry standard that enables tape to be used for backup, archive, and other applications without being tied to proprietary software.

- **Understand the advantages of tape** — for large repositories and data portability
- **Learn about LTFS-Open Edition and LTFS-Library Edition** — for single tape drives and libraries
- **Use LTFS-formatted tape** — for different industry uses and applications

Oracle engineers hardware and software to work together in the cloud and in your data center. For more information about Oracle (NYSE:ORCL), visit oracle.com.



Open the book and find:

- How explosive data growth is creating new digital archive challenges
- Why LTFS-formatted tape is the best mechanism for large file transport
- What traditional tape storage issues LTFS helps solve
- How different industries are using LTFS to address their archive needs

Go to Dummies.com
for videos, step-by-step examples,
how-to articles, or to shop!

FOR
DUMMIES[®]
A Wiley Brand

ISBN: 978-1-118-81066-8
Not for resale

These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.