



| IBM Almaden Research Center

Efficient Digital Media Workflows using the Linear Tape File System (LTFS)

MSST 4/18/2012

Arnon Amir, David Pease, Rainer Richter (IBM Almaden), Ed Childers (IBM Tucson)

Overview

- LTFS Introduction and Overview
 - The Digital Dilemma
 - The junction between Video and Data tapes
 - LTFS Single Drive and Library
- File-based Media Workflow Using LTFS
- Logical Copy and Physical Copy
- Future work



2011 Engineering Emmy® Award,
by the Academy Of Motion
Pictures, Arts and Sciences



IBM LTFS Product Debut, NAB 2010, Las Vegas



Hollywood Post Alliance Award® 2011



PICK HIT Award®,
NAB 2011, Las Vegas

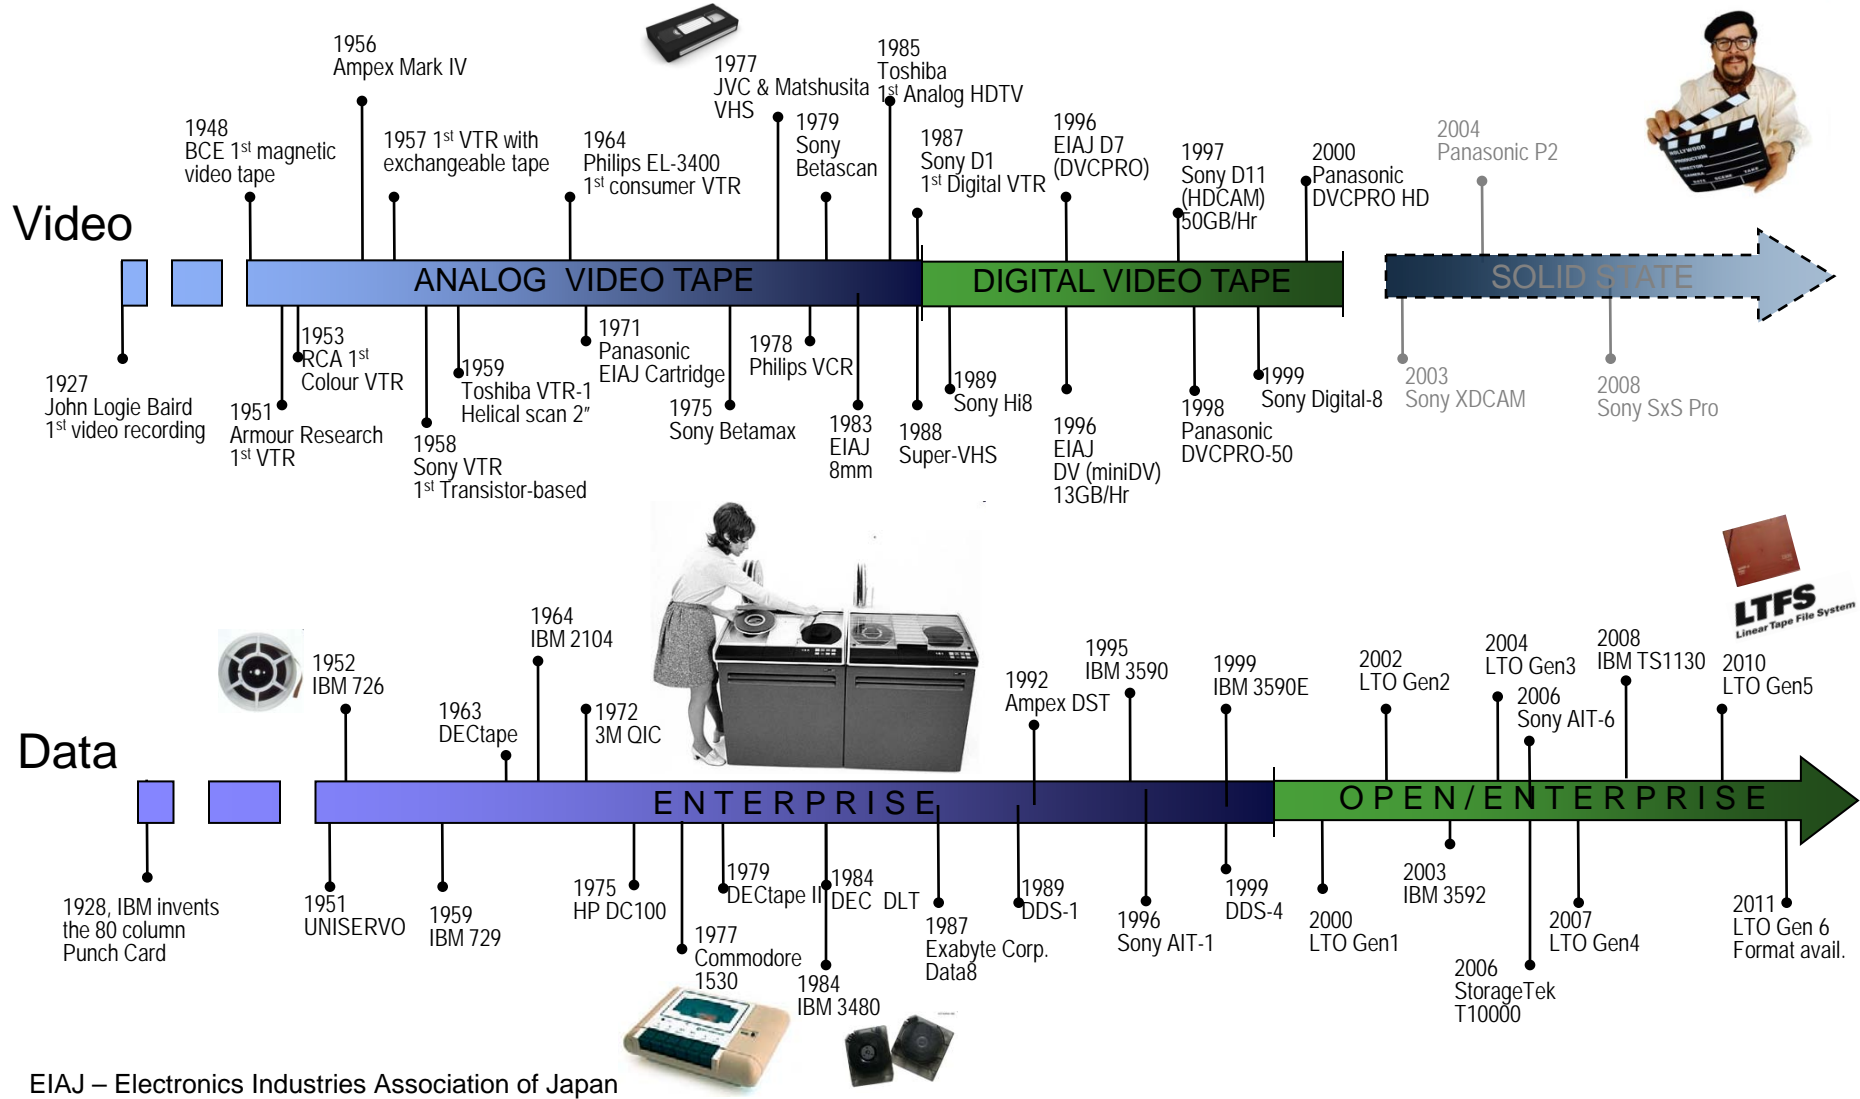
First IBM 726 Tape Unit (May 21, 1952)



- At 7,500 cps it was 56 times faster than punch cards
- Capacity in excess of 2 million digits a tape.
- With 100 bpi it could store the equivalent of 35,000 punch cards on a single, 1200 foot reel.
- 100% of the tapes were inspected with a microscope and knife...defects would be cut out and sent back to the manufacturer



60+ Years of Magnetic Data and Video Tapes (partial list)



The Digital Dilemma

[The S&T Council of the Academy (AMPAS), 2008]

The Digital Workflow Transformation: The Broadcast and Cinema industries change from video-tape to file-based workflows.

There are many benefits to file-based workflows.

However: "... a totally committed, binding switch to digital has one major drawback: the absence of guaranteed, long-term access to created moving image and sound content."

Requirement: Archive for 100 years with no loss.

- Manage future costs
- Minimize #migrations and disruption
- Minimize risk (e.g., single-vendor)

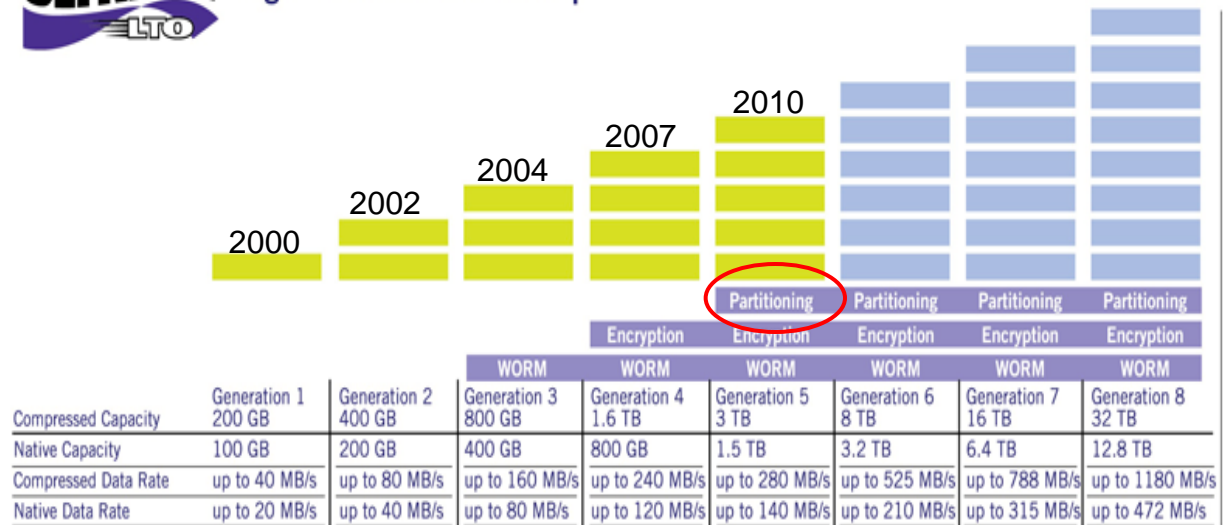


A Media Archive at NHK, Japan

LTO Tape & Roadmap

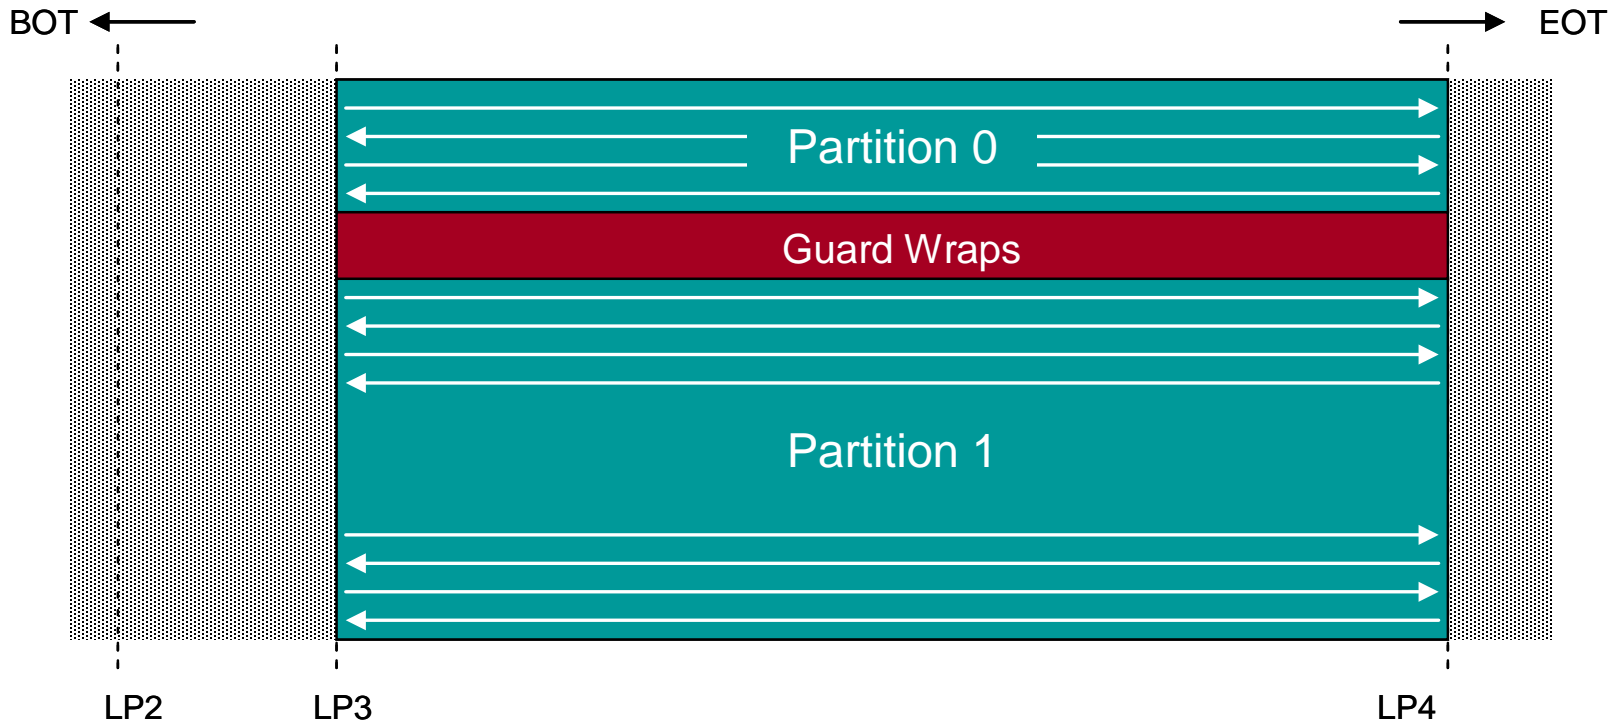
- ❑ Open standard by LTO Consortium: IBM, HP, Quantum.
- ❑ Dominates data tape market
- ❑ Gen-5 innovation of Dual Partitions enables LTFS.
- ❑ Backward compatibility
(Read: 2 Gen-s back, Write: 1 Gen. back)
- ❑ Forward Roadmap

Eight-Generation Roadmap



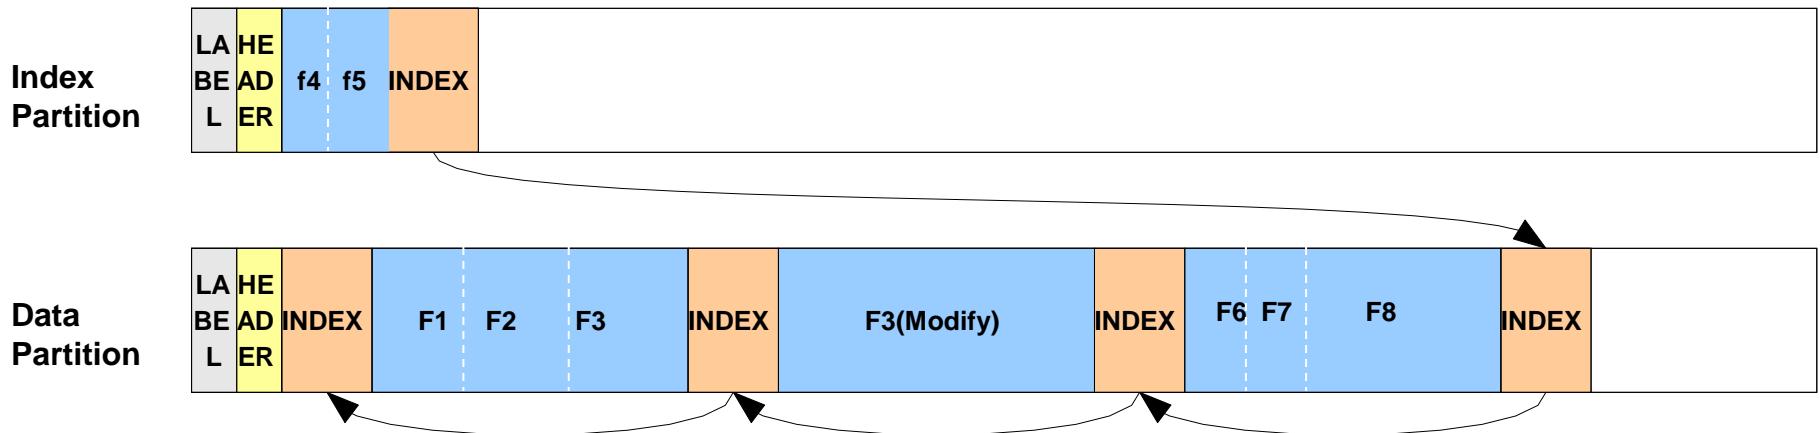
Note: Compressed capacities for generations 1-5 assume 2:1 compression. Compressed capacities for generations 6-8 assume 2.5:1 compression (achieved with larger compression history buffer).
 Source: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only.
 Linear Tape-Open, LTO, the LTO logo, Ultrium, and the Ultrium logo are registered trademarks of HP, IBM and Quantum in the US and other countries.

Dual-Partition Tape – A logical view



- ❑ Index partition: 2 wraps (about 37 GB)
- ❑ Data partition: 76 wraps (95% of 1.5 TB)
- ❑ Guard Wraps: 2 wraps (unused)

Index Arrangement on LTFS Tape



Tape contains File/Directory Index in Index Partition

Index kept as a human-readable XML schema

Keeps multiple “generations” (older versions of XML schema)

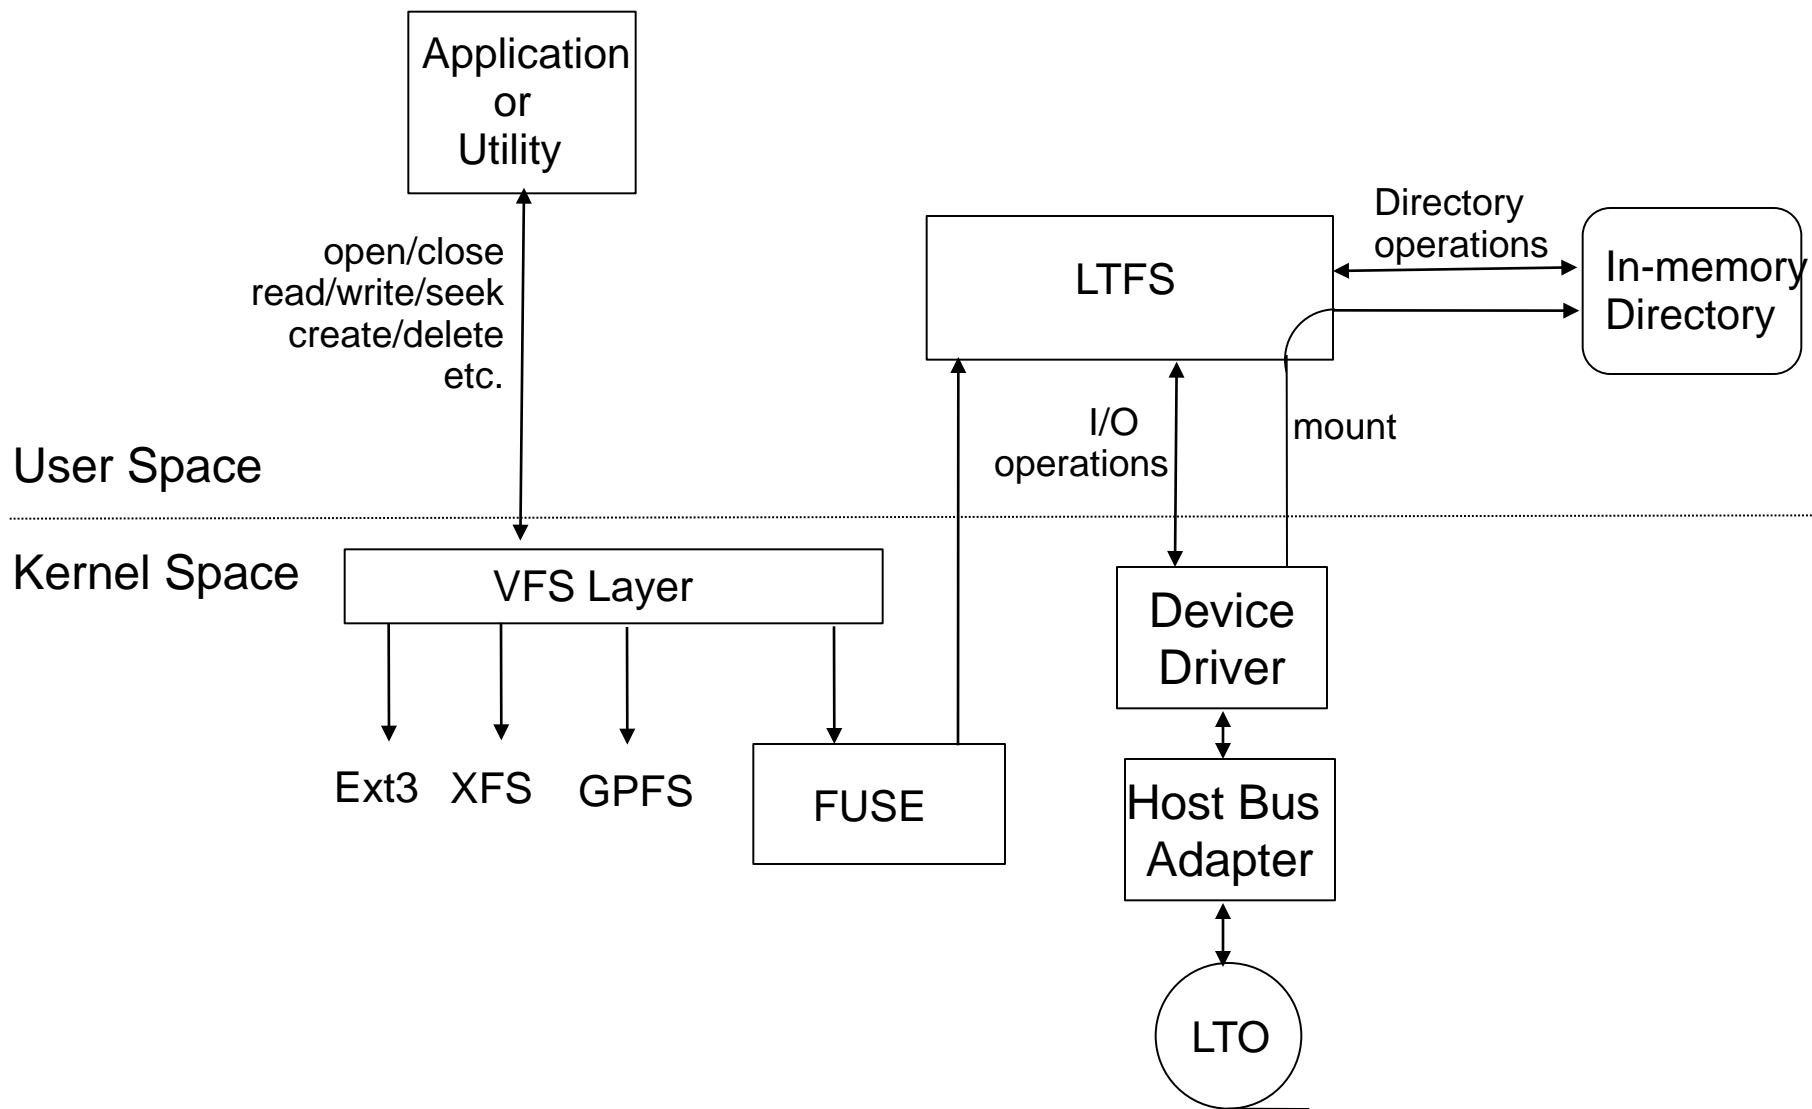
Links back to previous index tables. Supports rollback (at mount)

Automated detection and recovery from un-mount (*ltfsck*)

Supports user metadata as Extended Attributes (EAs)

LTFS Open format: http://www.trusttlo.com/LTFS_Format_To%20Print.pdf

LTFS Implementation Architecture (Linux)



LTFS Overview

- ❑ An open, non-proprietary tape format
 - ❑ Open source code and specification
 - ❑ Invented by IBM Almaden Research Center
- ❑ File system support for:
 - ❑ MAC, Linux, Windows.
- ❑ Multi-vendor support
 - ❑ Supported by the LTO Consortium
 - ❑ Media by 5 manufacturers
- ❑ Self-describing
 - ❑ Each cartridge has index of files on tape
- ❑ Mount and export as other file systems

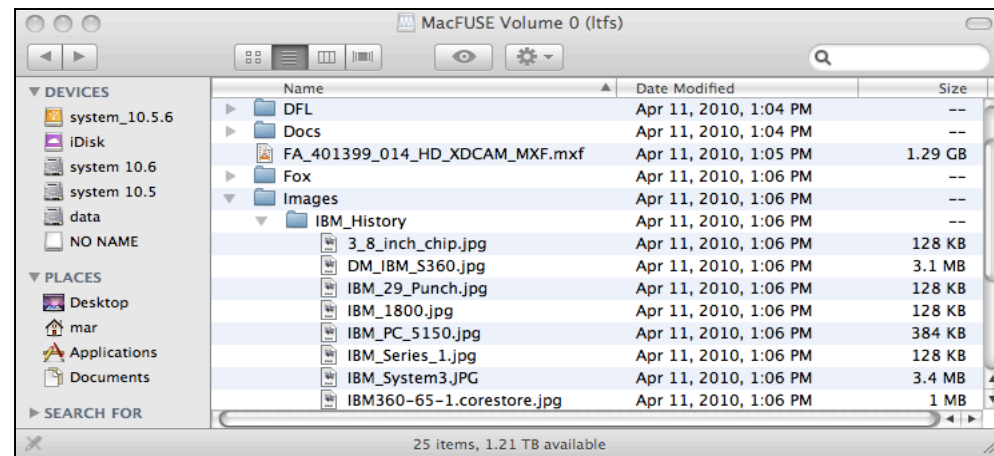
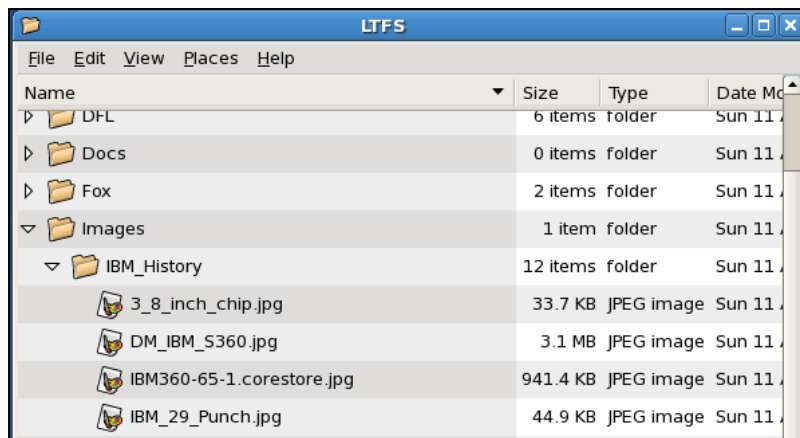


Enabling file based digital workflows and unlocking the digital tape value proposition for all phases of the media production, distribution, and archive life cycle

LTFS in Standalone Drive Mode

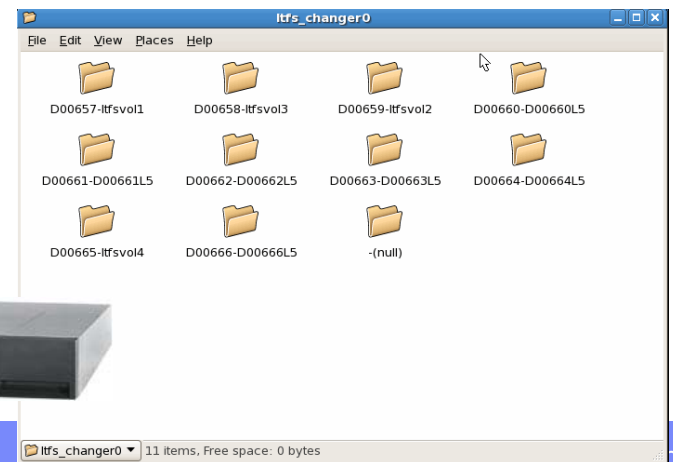
(Open Source: www.ibm.com/systems/storage/tape/lufs)

- ❑ Support at OS level for file commands (POSIX) – open, read, write, seek, ...
 - ❑ Enables standard applications to write/read LTFS tapes
 - ❑ Directories, Files, command-line and GUI file explorers
- ❑ Makes tape self-describing, portable, interchangeable, and long-term archive-able storage media
- ❑ For M&E:
 - ❑ One storage format for all video & multimedia data types
 - ❑ Self-describing: separate storage format from OS & storage system



LTFS in Library Mode

- ❑ Mount Library, not Drive
- ❑ LTFS presents directory structure of entire library
 - ❑ Each volume shows as separate file system directory
 - ❑ LTFS Caches Index in memory of each tape read/written
- ❑ After mount, all tape directories are viewable, searchable without mounting any tape
 - ❑ Meta-data operations (browse directory, filename search) do not require tape movement
 - ❑ LTFS drives automation to mount tape on file read/write on request
- ❑ Supported across entire automation platform

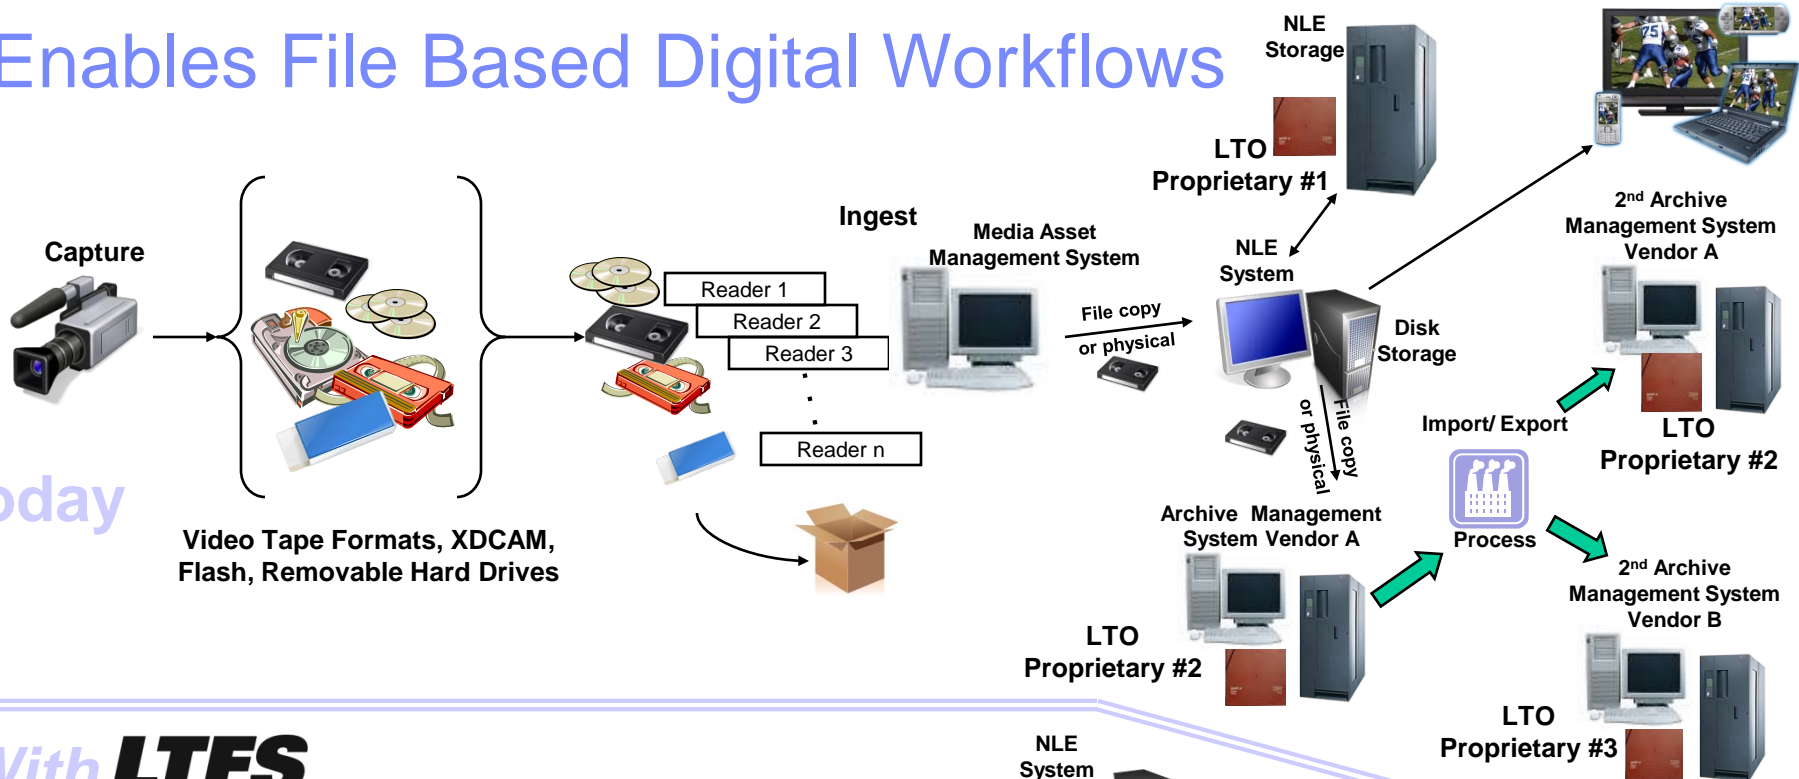


What LTO/LTFS is not ...

- ❑ The file system looks, to both users and applications, as if it's a Hard Drive. However seek time is still of a tape
 - ❑ Some applications might time-out while waiting for seek
 - ❑ For efficient use, user must be educated of the tape media
- ❑ LTO is an append only tape format.
 - ❑ LTFS allow to modify files by replacing, splitting and adding extents
 - ❑ To reclaim released space, entire tape need to be erased (active set copied)
- ❑ Better read one file at a time (can write multiple) otherwise might observe a lot of tape seeking...
- ❑ Beware of any indexing applications “peeking” at files (e.g., file search indexers, thumbnails generators, ...)

Enables File Based Digital Workflows

Today



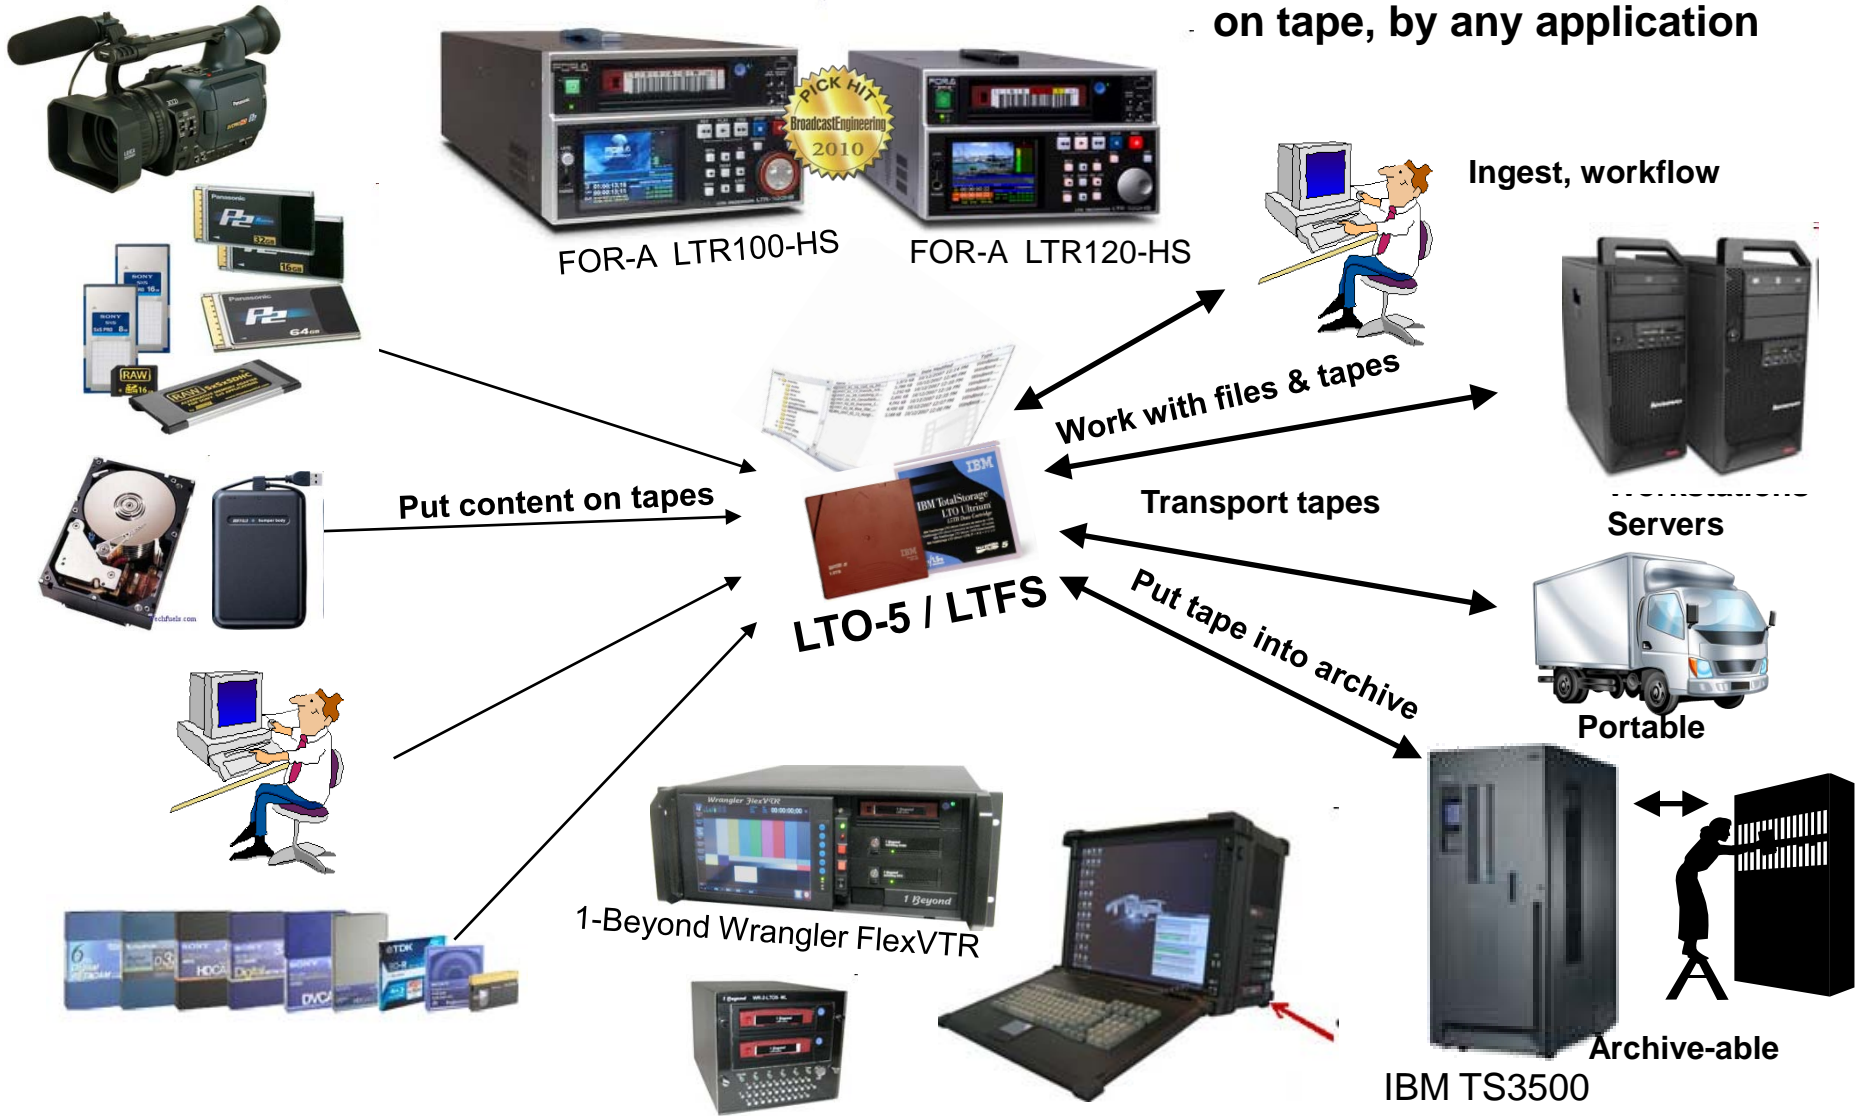
With **LTFS** Linear Tape File System



LTO5 (50 hours/cartridge)
D2M cassette (1.5 hours/cassette)
LTO Uses 1/200th vault space

Video Capture to LTO/LTFS

Access any file, directly on tape, by any application

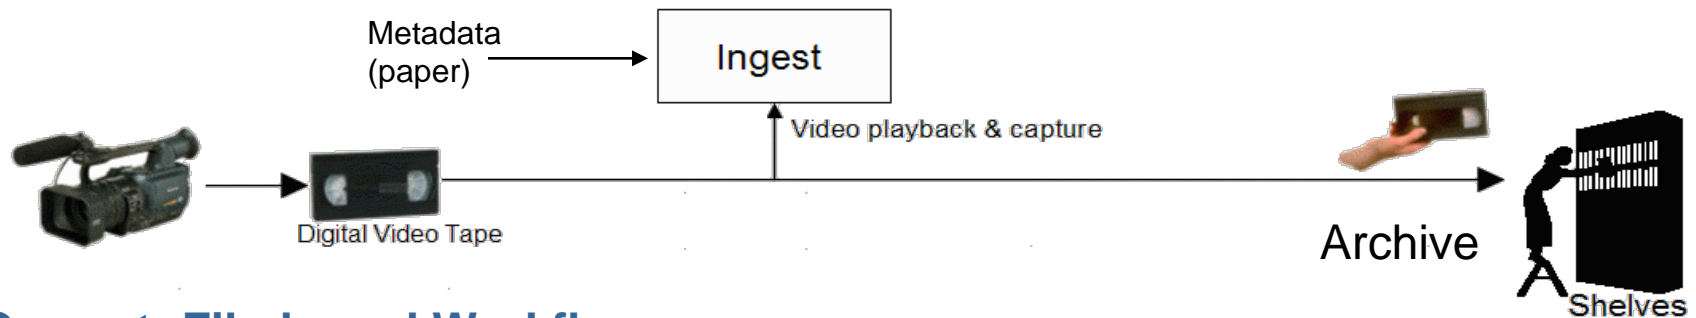


Media Workflow: From Cassette to File-Based

Before: Digital Video Cassette workflow

Pros: portable AND long term archive-able storage media


Cons: proprietary storage, single video format, very expensive, slow access, metadata on separate media.



Current: File-based Workflow

Pros: any FS storage, any video format, portable, easy access, metadata.

Cons: many FS storage media, portable is not archive-able, copy everything several times, proprietary archive format



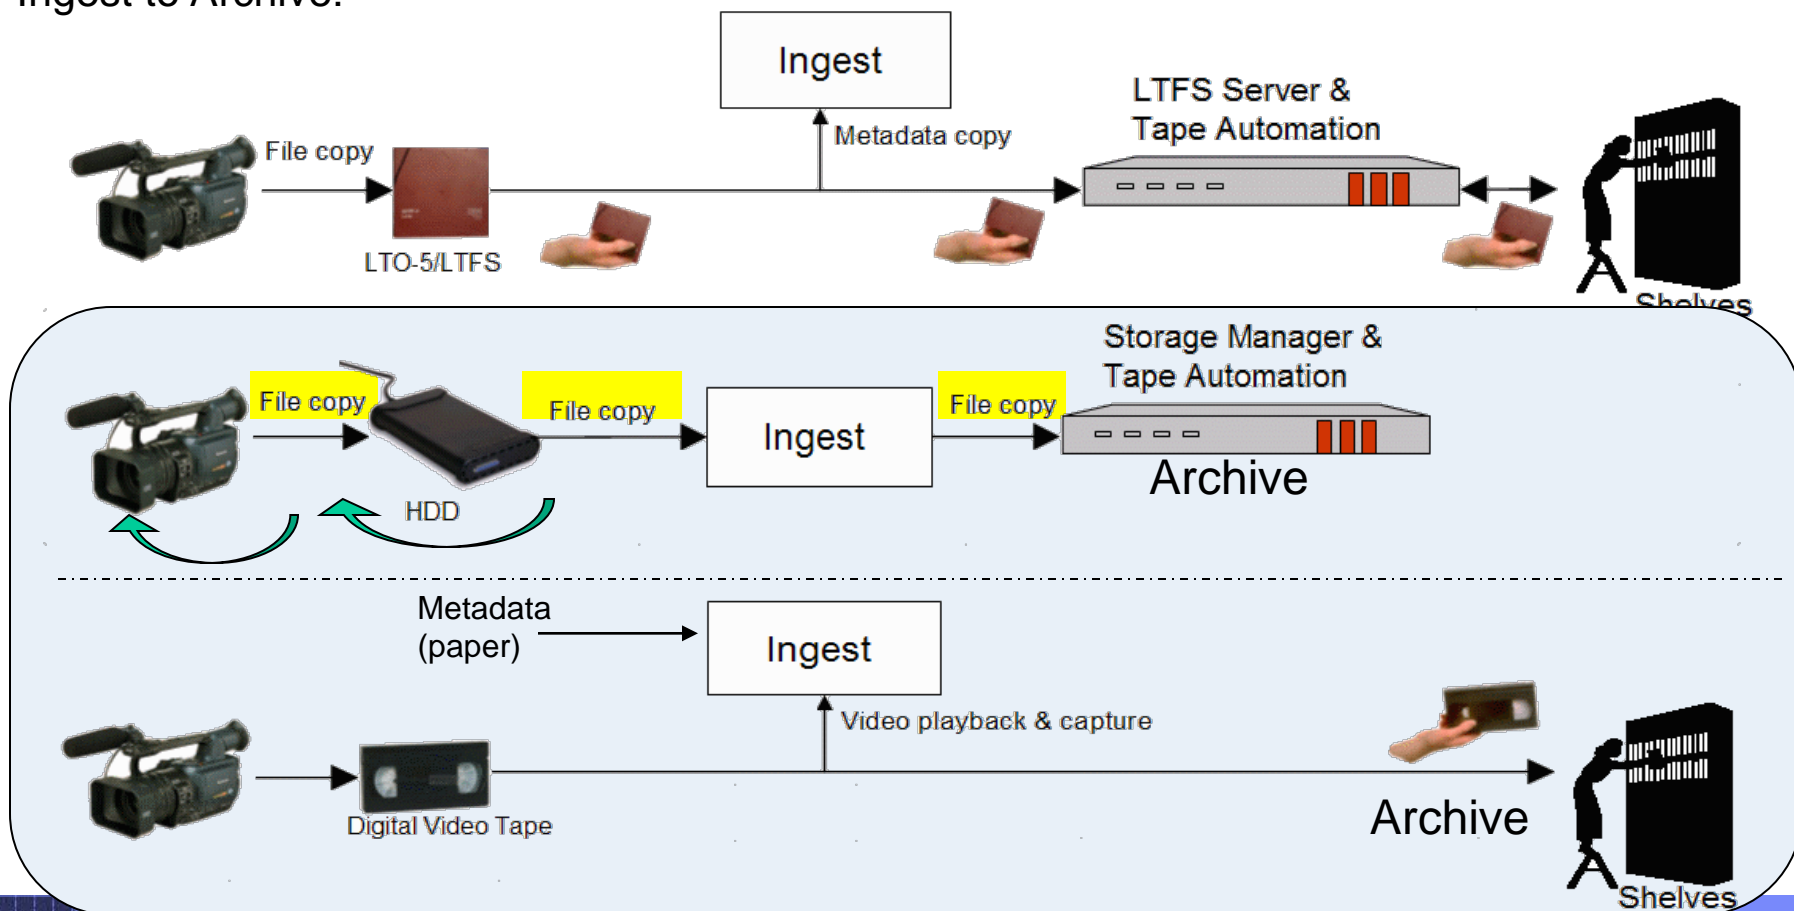
Recycle FLASH & portable storage media

File Based Workflow With LTFS:

Pros: open storage media, any video format, portable, metadata, archive-able, cost-effective (enjoy large-scale economy of data storage)

Cons: more IT-savvy than video tape, might cause long file-seek times (naïve use).

Ingest to Archive:

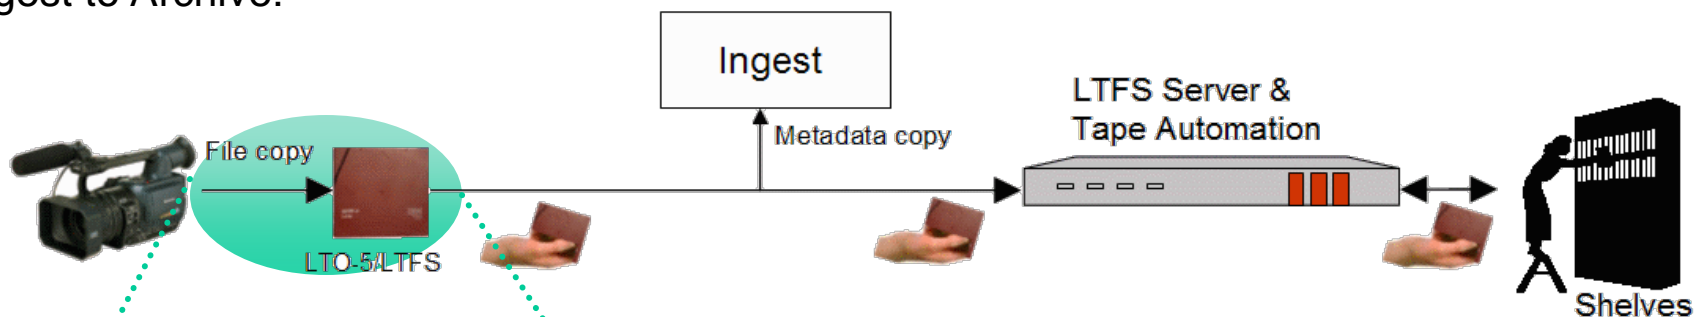


File Based Workflow With LTFS: MXF Metadata

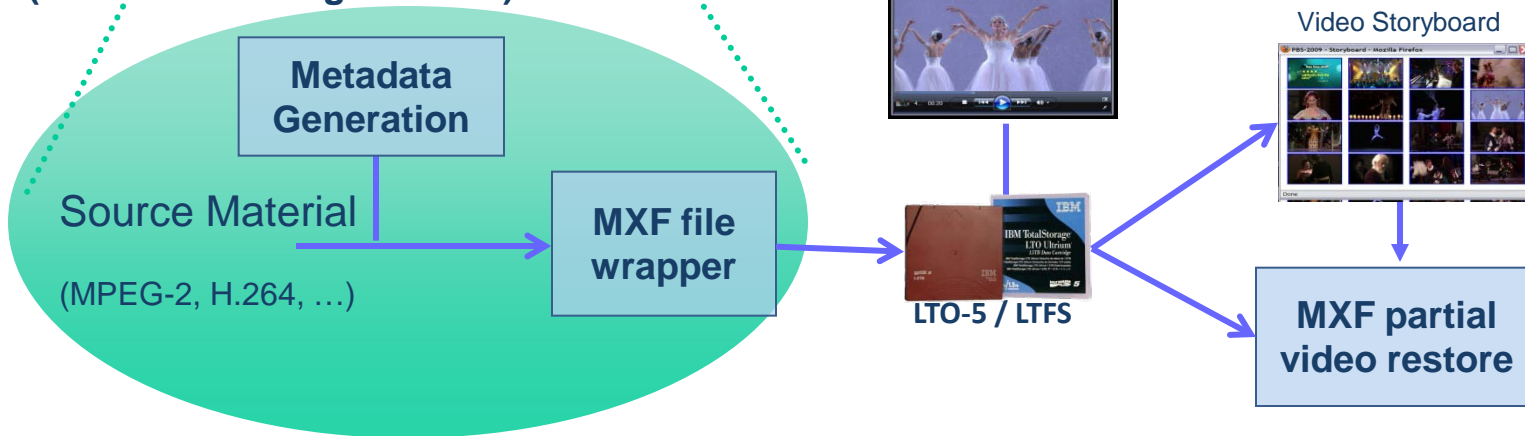
Pros: open storage media, any video format, portable, metadata, archive-able, cost-effective (enjoy large-scale economy of data storage)

Cons: more IT-savvy than video tape, might cause long file-seek times (naïve use).

Ingest to Archive:

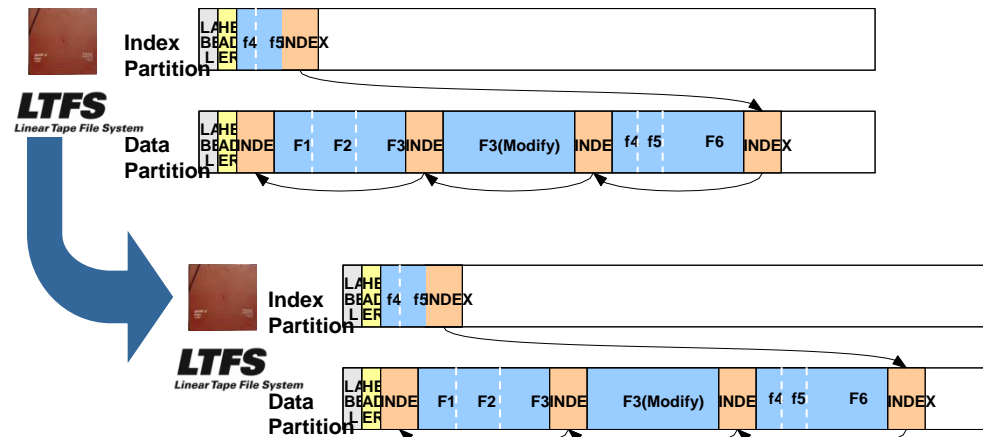
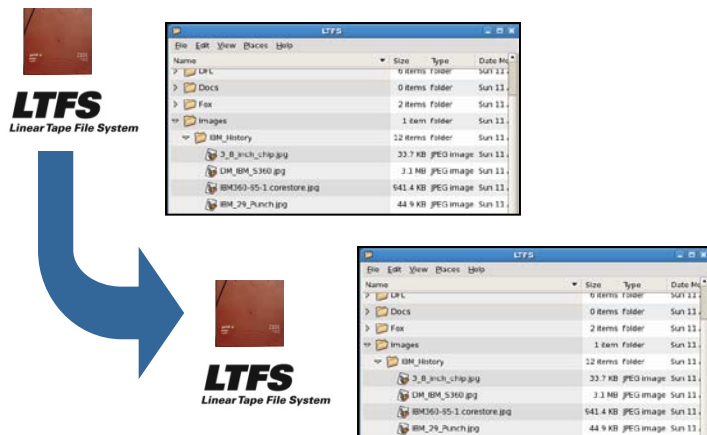


MXF Metadata (Material Exchange Format)



Logical Copy and Physical Copy

Logical Copy (IBM LTFS Copy Tool)	Physical Copy (IBM ITDT v6.1)
Copy files and directories (like cp) (tape->tape, disk->tape, tape->disk)	Duplicate tapes (like dd for LTFS) (tape->tape, image->tape, tape->image)
Optimal copy by files' order on tape	Copy blocks by sequential order
Preserve directories, metadata, EAs	Preserve partition, location
Copy all/selected files (console/batch)	Copy entire tape content
Copy current files, defrag	Preserve multiple extents
Reclaim tape space	Preserve deleted files & rollback



Logical Copy

- ❑ Operates at files level on LTFS-mounted tape/s

- ❑ Interactive mode: Files specified by regexp, for example:

```
ltfscp -srcpath=/mnt/LTFS/proj1 -dstpath=~ /proj1 -srcspec=*.dpx
```

- ❑ BATCH mode: Files (and options) specified by an XML script:

```
ltfscp -batchfile=~ /proj1 /batch21.xml
```

- ❑ Optimize copy performance:

- Reorder files list using files' order on tape to minimize total seek time

- ❑ Preserve file attributes, dates, and extended attributes

- ❑ Space reclamation (copy only current files), file de-fragmentation

- ❑ Multi-buffer, multi-file, multi-threaded copy (maximize tape streaming)

- ❑ Use cases:

- ❑ high performance batch jobs for archive and restore; tape reclamation, data migration

Logical Copy – batch mode

- ❑ Can specify any list of files
- ❑ Different levels of logs/reports
- ❑ Most suitable for system-generated batch jobs

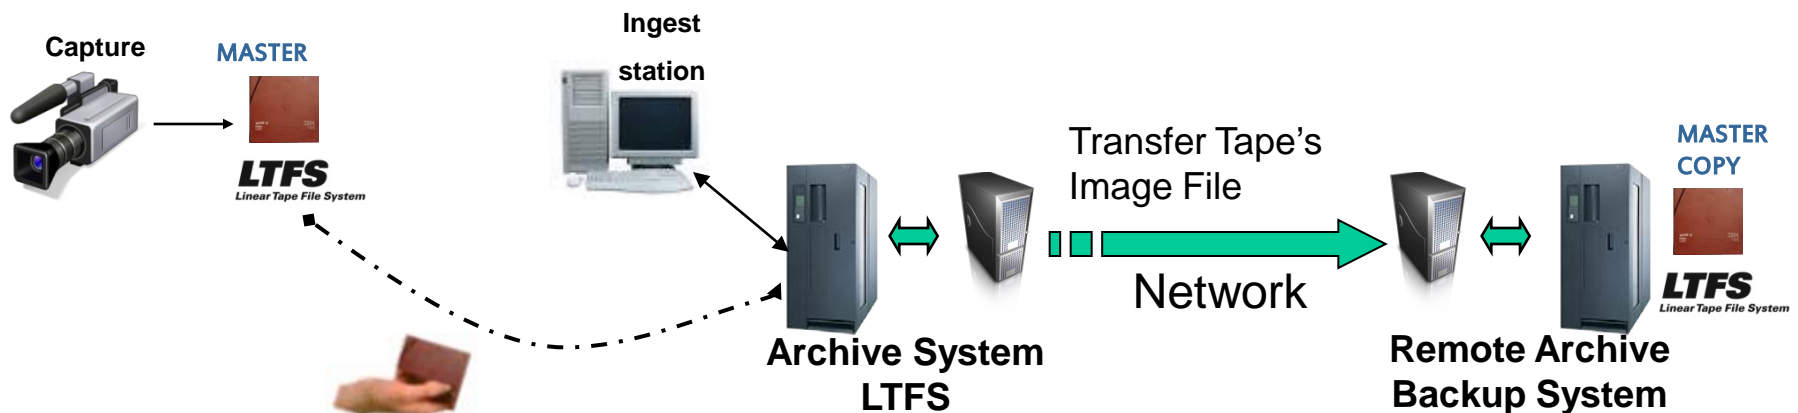
In this example:

- ❑ Traverse directories recursively
- ❑ Allow use of sparse files
- ❑ Copy two groups of files, one *.dpx and the second contains two specified files
- ❑ All warnings are logged in a log file

```
<?xml version="1.0" encoding="UTF-8"?>
<ltfscpspec version="1.0">
  <params>
    <loglevel>WARNING</loglevel>
    <recursive>enable</recursive>
    <sparse>enable</sparse>
  </params>
  <data>
    <file>
      <srcpath>/data/proj1</srcpath>
      <dstpath>/mnt/ltfs/proj1</dstpath>
      <srcspec>*.dpx</srcspec>
    </file>
    <file>
      <srcpath>/data/proj1</srcpath>
      <dstpath>/mnt/ltfs/proj1/doc</dstpath>
      <sf>copyright.doc</sf>
      <sf>readme.doc</sf>
    </file>
  </data>
</ltfscpspec>
```

Physical copy

- ❑ Copy all blocks sequentially by source partition & block order
- ❑ Multi-buffers, multi-threading architecture, maximize tape streaming
- ❑ While copying, adapt index blocks to comply with target tape
- ❑ Copy verification (optional)
- ❑ Handles and adapt LTFS information in CM (Cartridge Memory)
- ❑ Can create a disk image of an LTFS tape
 - ❑ Use in a single-drive mode
 - ❑ Use to copy tapes to remote location (archive mirror)



Case Study: Thought Equity Motion

Video Archiving in the Cloud

Challenges

- Needed a low cost delivery platform for enterprise scale Video Supply Chain as a Service
- Information growth of ~100 TB per month
- Easy self-serve access required by clients

Solution

- IBM Linear Tape File System at several global locations, including some client facilities
- IBM System Storage® TS3200 Tape Library, LTO®-5 tape drives

Benefits

- Opened up new business opportunities
- Enabled more predictable and transparent pricing for clients
- Portable, interoperable, scalable, cost-effective data protection and long-term storage




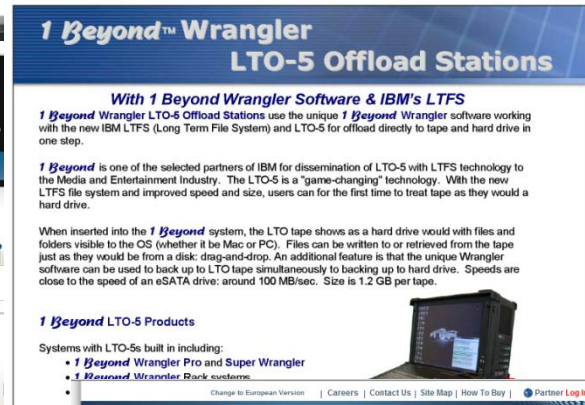

thoughtequity
MOTION.

*'LTO 5 and LTFS significantly reduce the ancillary costs around storage. This is a **real game-changer** from IBM'*

Mark Lemmons

CTO, Thought Equity Motion

LTFS Ecosystem



Challenges and Future Development

- ❑ LTFS Standardization
- ❑ Systems and Applications
 - ❑ Ingest/migrate and catalog.
 - ❑ transport, backup/restore, archive.
- ❑ Support large tape archives (100K-1M tapes)
 - ❑ Distributed across servers, vendors, geographies
 - ❑ Access via Hierarchical Storage Management (HSM)
- ❑ Domain-specific devices (e.g., video transcoding)

Acknowledgements

LTFS Research Team: Brian Biskeborn, Lucas Villa-Real, Michael Richmond, Anurag Sharma (Almaden) Atsushi Abe (IBM Yamato Lab)

LTFS Product Development Team, IBM STG (Tucson, Yamato & Mainz)

IBM Media & Entertainment Industry S&D.



Conferences April 14 – 19, 2012 // Exhibits April 16 – 19
Las Vegas Convention Center, Las Vegas, Nevada USA

Welcome, Guest. [\[Login\]](#)



myNAB Show	Exhibitors	Sessions	Download My Planner	Return to NABShow.com	
myNAB Show	My Profile	My Schedule	My Messages	My Exhibitors	My Sessions

[Go to EXHIBITOR Portal](#)

Other Resources

- [Register Now](#)
- [Exhibitor Search](#)
- [Session/Speaker Search](#)
- [Return to NABShow.com](#)
- [Need Help?](#)

Hot Session Topics

[Access Networks](#)
Streaming **Audio**
Telepresence Multiplatform
Workflow
Monetization
Broadband
Storage Transmission
[Advertising 3D](#)
Cameras **Mobile**
Sports
[Cinematography](#)
Social Media
[Production](#) Radio
Filmmaking Post

LTFS: Setting Your Videos and Data Free Now and for the Future!	
Session Date/Time	Wednesday, 4/18/2012 01:00 PM - 03:00 PM
Room Number	S224
Session Description	Learn how LTO-5 tape storage technology with the innovative Linear Tape File System (LTFS) can help save costs and make tape easier to use than ever before. Easily view tape content with directory tree structures, drag and drop files, and share video content across platforms. Come see why studio execs said "This is exactly what we need!" Also, hear from LTO solution providers; Crossroads, FOR-A, Panasonic, SGL, Storage DNA and XenData on how they use LTO and LTFS technology to help make your video storage workflow and archive easier and cost effective!
Speakers	Laura Loreda, Ultrium LTO Bruce Master, Ultrium LTO Mark Pastor, Ultrium LTO David Cerf, Crossroads Pedro Silvestre, FOR-A Phil Storey, XenData, Inc. Howard Twine, Software Generation Limited Ryota Yasutomi, Panasonic
Session Sponsor Logo	