

Technical Support Services Contract (TSSC)				
USGS EROS	G15PC00012	Deliverable ID	007-1.110.1.W01.1.2.6-A4703	
FY16 Offline Archive Media Trade Study		Date	07/29/16	Version Nbr 1.0

ARCHIVE AND RECORDS MANAGEMENT

FISCAL YEAR 2016 ARCHIVE MEDIA TRADE STUDY



TSSC FISCAL YEAR 2016 OFFLINE ARCHIVE MEDIA TRADE STUDY

Remote Sensing, CEOS, and Archives Coordination Project

By Tom Bodoh¹

Preface

This document contains the Offline Archive Media Trade Study prepared by Stinger Ghaffarian Technologies, Inc. for the U.S. Geological Survey. This trade study presents the background, technical assessment, test results, and recommendations.

The U.S. Geological Survey uses trade studies and reviews for internal purposes and does not endorse vendors or products. The results of the study were determined by criteria weights selected by the U.S. Geological Survey to meet its unique requirements. Other organizations could produce different results by altering the criteria weights to meet their own requirements.

Acknowledgement

The author would like to thank the following staff who provided input and review: Brent Nelson, Ken Gacke, Dave Nelson, and Brian Reutzel.

¹ Stinger Ghaffarian Technologies, Inc., contractor to the U.S. Geological Survey, work performed under U.S. Geological Survey contract G15PC00012.

DCN: 007-1.110.1.W01.1.2.6-A4703

Version: 1.0

Contract: G15PC00012

Date Modified: 7/29/16

Notice: Paper copies of this document are uncontrolled. A controlled electronic version is located in TSSC Task Management Document Repository.

Abstract

This document is a trade study comparing offline digital archive storage technologies. The document compares and assesses several technologies and recommends which technologies could be deployed as the next generation standard for the U.S. Geological Survey Earth Resources Observation and Science Center. Archives must regularly migrate to the next generation of digital archive technology, and the technology selected must maintain data integrity until the next migration. This document is the fiscal year 2016 revision of a study completed in Fiscal Year 2001 and revised in Fiscal Years 2003, 2004, 2006, 2008, 2010, 2012, and 2014.

Contents

Preface	2
Acknowledgement.....	2
Abstract	3
Abbreviations	8
Revision History	10
February 2004.....	10
September 2006.....	10
June 2008	11
June 2010	11
June 2012	11
August 2014.....	12
July 2016.....	12
Introduction.....	13
1.1 Purpose and Scope.....	13
1.2 Background	14
1.3 Data Integrity.....	15
1.4 Selection Criteria.....	17
1.5 Dismissed Technologies	18
1.5.1 Magnetic Disk	19
1.5.2 Solid State Disk (SSD).....	19

1.5.3 DVD, Blu-Ray	19
1.5.4 Newer Storage Technologies.....	20
1.5.5 Cloud Storage.....	20
Technical Assessment.....	22
2.1 Analysis.....	22
Oracle T10000D.....	24
Advantages	24
Disadvantages.....	25
Summary	25
Oracle T10000E	27
Advantages	27
Disadvantages.....	28
Summary	29
IBM LTO7.....	30
Advantages	30
Disadvantages.....	31
Summary	31
IBM LTO8.....	33
Advantages	33
Disadvantages.....	33
Summary	34
IBM TS1150	35
Advantages	35

Disadvantages.....	35
Summary.....	36
IBM TS1160.....	37
Advantages.....	37
Disadvantages.....	37
Summary.....	37
Tables.....	39
3.1 Design criteria.....	39
3.2 Transfer Rate.....	40
3.3 Capacity.....	41
3.4 Cost Analysis.....	42
3.5 Scenarios.....	43
3.6 Vendor Analyses.....	44
3.7 Drive Compatibility.....	45
3.8 Ranking Summary.....	45
Conclusions and Recommendations for USGS Offline Archiving	
Requirements.....	47
4.1 Weighted Decision Matrix.....	47
4.2 Conclusions and Notes.....	49
4.3 Recommendations.....	50
References Cited.....	52
Vendor Sites.....	53

Consortium Sites 53

Other 53

Figures

Figure 1. Oracle Roadmap (uncompressed) (Oracle, 2015). 26

Figure 2. LTO Roadmap (compressed) (source: LTO Consortium). 32

Figure 3. IBM Roadmap (uncompressed) (Source: IBM) 36

Tables

Table 1. Recent and Current Offline Archive Technologies Used at The Earth Resources Observation and Science Center (current in bold) 15

Table 2. Tape Drive Markets and Characteristics 17

Table 3. Technology Comparison (yellow-highlighted columns indicates unverified information) 23

Table 4. Design Criteria and Target Market 40

Table 5. Transfer Rates 41

Table 6. Storage Capacities 41

Table 7. Drive and Media Costs 42

Table 8. Scenario Costs (drives, media) 43

Table 9. Vendor Analyses 44

Table 10. Drive Compatibility 45

Table 11. Ranking Summary (blue indicates the highest ranking in category) 46

Table 12. Weighted Decision Matrix 48

Abbreviations

BaFe	Barium Ferrite
BER	Bit Error Rate
CD-ROM	Compact Disc – Read Only Memory
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
DCT	Digital Cassette Tape
\$/unit	dollars per unit
\$/TB	dollars per terabyte
DLT	Digital Linear Tape
DVD	Digital Video Disc
EO	Erasable Optical
EROS	Earth Resources Observation and Science
est	estimated
FYyy	Fiscal Year yy
GB	Gigabytes (1,024 MB, or 1,073,741,824 bytes)
Gbit	Gigabits (1,073,741,824 bits)
HD-DVD	High Definition-Digital Versatile Disc
HDT	High Density Tape
HP	Hewlett-PackardHSM Hierarchical Storage Management
HVD	Holographic Versatile Disc

HW	Hardware
I/O	Input/Output
IBM	International Business Machines
LTO	Linear Tape Open
LTFS	Linear Tape File System
MB	Megabytes (1,048,576 bytes)
MP	metal particle
NARA	National Archives and Records Administration
%	percent
Q1, Q2, Q3, Q4	Fiscal or calendar quarter
QIC	Quarter-Inch Cartridge
SAIT	Super Advanced Intelligent Tape
sec	Second
SGT	Stinger Ghaffarian Technologies, Inc.
SSD	Solid State Disk
STK	StorageTek (subsequently bought by Sun, which was bought by Oracle)
TB	Terabytes (1,024 GB or 1,099,511,627,776 bytes)
TBD	To be decided/determined
USGS	United States Geological Survey

Revision History

February 2004

- Added revision history page. No revision history is available for the FY03 revision.
- Changed to allow for consideration of helical scan if certain performance criteria are met.
- Added LTO2 as a current archive technology.
- Added SAIT-1 and SDLT 600 as considered drives.
- Replaced IBM 3590 with IBM 3592.
- Removed LTO1 and SDLT 320 from the study.
- Considered all drives in the study.
- Increased the minimum specifications for capacity and transfer rate.
- Reworked cost scenarios and reduced the number of cost scenarios to three.
- Removed transfer-time scenarios.
- Removed maintenance from cost scenarios.
- Removed criteria indicating multi-vendor availability as an advantage.

September 2006

- Overall refresh of study.
- Revised description of drive classes (enterprise, backup).
- Added LTO3, TS1120, T10000, and DLT-S4 as current technologies and removed drives they replaced.
- Added LTO4 and SAIT2 as future technologies.
- Made vendor analyses formula more equitable by increasing weighting of company age.

- Added citation appendix.

June 2008

- Overall refresh of study, removing most references to older technologies.
- Added disk as a dismissed technology.
- Changed LTO4 to a current technology.
- Added T10000B, LTO5, and TS1130 as future technologies; deleted LTO3, SAIT1, and SAIT2.
- Modified so that future technologies are no longer scored.
- Decreased the number of drives for scenarios 2 and 3.

June 2010

- Overall refresh of study, removing most references to older technologies (T10000, LTO4, DLT).
- Changed T10000B, LTO5, and TS1130 to current technologies.
- Added T10000C, LTO6, and TS1140 as future technologies.
- Removed maintenance costs because of lack of data.
- Adjusted minimum transfer rate and capacity to be considered for the study.

June 2012

- Overall refresh of study, removing most references to older technologies (T10000B, TS1130)
- Changed T10000C and TS1140 to current technologies.
- Added T10000D and TS1150 as future technologies.
- Removed references to CD-ROM, DLT 8000, QIC, Mammoth, Erasable Optical (EO), HD-DVD, and 9840 under dismissed technologies.
- Removed row from table that showed all drives use the same offline storage shelving.
- LTO drive price is now for robotic drives.

- Removed future drives from analysis tables.
- Removed drive warranty row from table.

August 2014

- Overall refresh of study, removing most references to older technologies (T10000C, LTO5)
- Removed most mention of DLT.
- Added mention of cloud.
- Changed T10000D and LTO6 to current technologies.
- Added T10000E and LTO7 as future technologies.
- Increased the minimum criteria for the study to a capacity of 2 TB and a transfer rate of 150 MB/sec.
- Increased the number of passes in the design criteria formula to accommodate increasing pass specifications.
- Tape drive cost estimate now includes 1-yr warranty/support.
- Increased the total capacity for each of the three scenarios.
- Changed the drive compatibility table and calculation to improve applicability to archiving.

July 2016

- Overall refresh of study, removing most references to older technologies.
- Changed LTO7 and TS1150 to current technologies.
- Added LTO8 and TS1160 as future technologies.
- Removed reference to Sony SAIT and Tandberg VXA-320 which are retired.

Introduction

1.1 Purpose and Scope

Typically, the purpose of a trade study is to analyze several courses of action and to provide the necessary information for the sponsor to reach a conclusion. In other cases, a trade study may revalidate an ongoing course of action.

This document assesses the options for the next generation of offline digital archive storage technology to be used for the digital archives of the U.S. Geological Survey (USGS). The selected technology must be capable of safely retaining data until space, cost, and performance considerations drive the next media migration. Data must be migrated before integrity degrades.

Nearly all of the USGS working archive holdings now reside on nearline robotic tape storage and are backed by an offline master copy. The nearline copy is referred to as the working copy. An ongoing need exists for offline storage for infrequently used working copies and for master and offsite copies where the working copy is stored nearline. An offline copy stored in a secure offsite location reduces the chances of corruption or tampering; online or nearline methods are susceptible to intentional or unintentional corruption, no matter the probability.

Linear Tape Open (LTO) has been the offline archive media of choice at the USGS Earth Resources Observation and Science (EROS) Center since 2003. There is no compelling reason for the USGS to change technologies at this time, and given the advantages of intergeneration read compatibility in an offline archive environment, there will be a continued interest in “staying the course” with LTO technology for the foreseeable future.

This predisposition to use LTO technology does not negate the need to periodically revisit offline storage technologies to stay informed of changes. When, or if, LTO no longer meets EROS requirements, this study (in future revisions) will show the way to the emerging replacement.

This study specifically does not address the online and nearline technologies used at EROS. The primary nearline mass-storage system at EROS contains a Hierarchical Storage Management (HSM) system using an Oracle SL8500 robotic tape library, Oracle T10000D tape drives, Oracle LTO7 tape drives (capable of reading LTO5 and LTO6 media), an Oracle host server, Oracle Hierarchical Storage Management (HSM) software, and a multi-vendor disk cache. A trade study using a different set of requirements than this study determined the architecture of this HSM.

This study determines the best offline archive media to meet EROS requirements. The findings of this study should not be misconstrued as an analysis of any specific technology for other purposes, such as enterprise backup or robotic nearline storage. Changing the criteria weighting factors would produce different findings tailored to other specific circumstances.

1.2 Background

The USGS EROS Center, near Sioux Falls, South Dakota, has archived offline datasets using several technologies. Table 1 shows the offline archive tape media used at EROS since tape archiving began, with the currently used media shown in bold.

Table 1. Recent and Current Offline Archive Technologies Used at The Earth Resources Observation and Science Center (current in bold).

Tape drive technology	Years used at EROS	Capacity	Transfer rate	Type
HDT	1978–2008	3.4 GB	10.6 MB/sec	Analog
3480	1990–2003	200 MB	2.0 MB/sec	Digital
DCT (Ampex DCRsI)	1992–2007	45 GB	12.0 MB/sec	Analog
3490	1995–2003	900 MB	2.7 MB/sec	Digital
DLT 7000	1996–2006	35 GB	5.0 MB/sec	Digital
SuperDLT 220	1998–2008	110 GB	10.0 MB/sec	Digital
Oracle 9940B	2002–2011	200 GB	30.0 MB/sec	Digital
HP LTO Ultrium 2	2003–present	200 GB	40.0 MB/sec	Digital
HP LTO Ultrium 3	2005–present	400 GB	80.0 MB/sec	Digital
HP LTO Ultrium 4	2007–present	800 GB	120.0 MB/sec	Digital
HP LTO Ultrium 5	2010–present	1.5 TB	140.0 MB/sec	Digital
Oracle T10000C	2012–present	5.0 TB	240.0 MB/sec	Digital
Oracle T10000D	2014–present	8.5 TB	252.0 MB/sec	Digital
HP LTO Ultrium 6	2013–present	2.5 TB	160.0 MB/sec	Digital
IBM LTO Ultrium 7	2016–present	6.0 TB	300.0 MB/sec	Digital

As technology advances, datasets grow and media ages, and as USGS Digital Library space fills, the USGS must migrate data to newer, more cost-effective, more physically compact, and higher performing storage technologies.

1.3 Data Integrity

Because the foremost goal of an archive is data preservation, data integrity must be the primary criterion for the selection of the drive technology. The following listed elements contribute to data integrity:

- The number of archival copies — USGS archives must have working and master copies, and an offsite copy is desirable. The master and working copies would ideally utilize different media types so that media or drive issues do not risk both copies.
- Drive reliability — a slightly less reliable drive technology can be used, but only with a sufficient number of copies in the archive.

- The storage location and environment — storage location and environment are a constant for all the technologies assessed because all EROS media are stored in a secure and climate-controlled environment.
- The composition of the media — some media compositions last substantially longer than others, but all the technologies in this study use similar long-lasting media compositions.
- Tape handling within the drive — this characteristic defines how a tape is handled by the drive: whether contact is made with the recording surface, how many serpentine passes are required to read or write an entire tape, and the complexity of the tape path.
- Error handling — Drives typically minimize data loss through Cyclic Redundancy Check (CRC) or other data recovery methods, and allow data to be read after skipping past an error. Though error detection on write is required, additional attention to data recovery on read is a higher priority because media degradation will eventually lead to read errors.
- Primary market — This criterion describes the target market of a drive and the characteristics of drives in that market:
 - A drive targeted to the backup market is designed for write many/read rarely and depends more on write error detection because the data are still available and can be easily rewritten. Backup drives are typically built for speed, capacity, and low cost.
 - A drive targeted to the enterprise market is designed for write many/read many use in a robotic library or auto-stacker, and equal emphasis is placed on detecting errors on read and write. Enterprise drives are typically built for reliability and speed, with capacity a secondary factor. Cost is a not a primary consideration.
 - A drive targeted to the archival market would be designed for write once/read rarely, and equal emphasis would be placed on detecting errors on read and write; however, no

drives are currently designed or marketed primarily for archiving. Most vendors would argue that their products are archive devices, but if forced to choose their primary market, no vendor would choose the limited archive market over the lucrative backup or enterprise markets.

Table 2. Tape Drive Markets and Characteristics.

Primary market	Reliability	Usage	Driving design factors
Backup	Moderate	Write many, read rarely	Low cost, high capacity, high speed
Enterprise	High	Write many, read many	High duty cycle for drives and media used with robotics
Archive	High	Write once, read rarely	Long-term reliability

The reliability of a long-term archive technology relates primarily to the long-term viability of the recorded media. Reliability in technology is difficult to determine except in retrospect because a technology needs to be implemented early enough in the life cycle so that drives can be kept working during the lifetime of a given media (or replaced with newer backward-compatible models). This study bases the reliability assessment on past experience with the vendor and their products, on specifications, on the experiences of others, or on experience gained from benchmarking.

1.4 Selection Criteria

The following criteria were used to determine which offline technologies should be considered:

1. The technology must be currently available and the most recent generation in a technology lineage. Drives that are anticipated/announced but not available are mentioned but not ranked in the final analysis.
2. The technology must have a capacity of at least 4 terabytes (TB) [4,000 gigabytes (GB)] of uncompressed data.

3. The technology must have an uncompressed write transfer rate of at least 200 megabytes per second (MB/sec).
4. The technology must use media that can remain readable for at least 10 years in a controlled environment. The lifetime of 10 years was selected because 10 years is the longest that a media technology would conceivably be used before space and transfer rate concerns would dictate a move to a new technology. Maintaining obsolete drives also becomes difficult and expensive after 10 years.
5. The technology must not be hampered by a poor reliability or performance history; for example, helical scan technologies such as 4 millimeter (mm), 8 mm, Digital Audio Tape (DAT), and D3 have proven to be unreliable in the past.

The following currently available drive technologies were selected for consideration:

1. Oracle T10000D
2. Oracle branded International Business Machine (IBM) LTO7 (Linear Tape Open)—representative of models by Quantum and Hewlett Packard (HP).
3. IBM TS1150

The following future drive technologies are mentioned but not considered:

1. Oracle T10000E
2. Oracle branded IBM LTO8
3. IBM TS1160

1.5 Dismissed Technologies

The following technologies were dismissed from analysis or consideration.

1.5.1 Magnetic Disk

Disk prices continue to drop, whereas reliability, performance, and capacity increase. Cost, management overhead, cooling, and power are considerations in using disk technology to archive large datasets. In the past several years it has become feasible to store the working copy of some datasets, or parts of datasets, on disk as long as archive copies are retained, typically on tape. Although tape media could remain viable for as many as 10 years, the more costly disk typically is replaced every 4 or 5 years to maintain supportability, reliability, space density, and performance. Serving frequently used working copies on disk provides significant performance benefits, although an offline master copy must be retained. Disk is not designed, or often used, for offline storage.

1.5.2 Solid State Disk (SSD)

Similar to magnetic disk, SSD prices continue to drop, whereas reliability, performance, and capacity increase. It is expected that SSD, with time, will replace magnetic disk for online storage. SSD does offer some benefits regarding archive storage—it is expected to tolerate long shelf storage better than magnetic disk, which suffers from coating deterioration. Even though SSD could become an option for future offline archive storage, it is too expensive to compete at this time and is not intended for offline storage.

1.5.3 DVD, Blu-Ray

Digital Video Disc (DVD) and related technologies seem promising from the standpoint of expected longevity of the media; however, optical media can degrade and become unusable in as little as 5–10 years (OSTA, 2003). Low capacity per media, low transfer rates, lack of media protection (no shell), no single standard, and high media costs describe a product that simply will not work for high-volume archival use.

Blu-Ray would certainly have some application in distribution and short-term storage of large amounts of data, but like the CD and DVD, Blu-Ray suffers from high media costs and low transfer rates, and given optical media history, the shelf longevity must be proven before being trusted in an archive environment. A follow-on higher capacity Blu-Ray is anticipated, though it will not meet the evaluation criteria for this study. The market for optical disc is driven by entertainment but is dwindling due to availability of online content.

1.5.4 Newer Storage Technologies

Several high-capacity optical disk technologies have been in the development phase for the past few years. Of the technology proposals that have appeared in trade journals and at conferences, none are available.

One high-tech example of future technologies is holographic storage. Products have been repeatedly announced, but have yet to ship. Holographic Versatile Disc (HVD) specifications indicate a planned capacity of 3.9 TB per disk and a transfer rate of 125 MB/sec (SearchStorage, n.d.).

Another example of potential future developments is a recent announcement by Sony of a magnetic substrate technology which could result in a tape product with a capacity of up to 185 TB (Sony, 2014). There is no indication if, or when, a tape drive will become available, or if it would be a Sony product.

1.5.5 Cloud Storage

Though not an offline media, cloud storage is emerging as a viable offsite storage alternative which could be used as one copy of an archive. At the present time, online public-cloud storage would be prohibitively expensive for rarely accessed deep archive of petabyte-scale datasets, but could eventually be leveraged as a working copy of limited datasets as prices continue to drop. Deep-archive

public-cloud storage would still utilize tape and would likely cost more than storing tapes at the National Archives and Records Administration (NARA). Cloud storage would not eliminate the mandate for deep archive storage at NARA, though NARA may also evolve its archive services to include a private cloud.

Technical Assessment

2.1 Analysis

This technical assessment includes drives selected for final evaluation (T10000D, LTO7, and TS1150) and drives anticipated to be released in the next two years (T10000E, LTO8, and TS1160). LTO drives are available from multiple vendors (Quantum, IBM, and HP), with an Oracle branded IBM drive selected to represent LTO technology in this study. The following tape technologies will be evaluated, but only the drives shown in bold will be included in the analysis and final evaluation:

- **Oracle T10000D**
- Oracle T10000E
- **IBM LTO7**
- IBM LTO8
- **IBM TS1150**
- IBM TS1160

Table 3. Technology Comparison (yellow-highlighted columns indicates unverified information)

Specification	T10000D	T10000E	IBM LTO7	IBM LTO8	TS1150	TS1160
Uncompressed capacity	8.5 TB ²	12–16 TB	6.0 TB	12.8 TB	10 TB	20 TB
Uncompressed xfer rate	252 MB/sec	400–600 MB/sec	300 MB/sec	472 MB/sec	360 MB/sec	720 MB/sec
Recording technology	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine
Tracks	4,608		3,584		5,120	
Channels	32	32 or 64	32	32	32	32 or 64
Passes ³	144		112		160	
Tape velocity (read)	4.75 m/sec		6.8 m/sec est		6.22 m/sec est	
Type	Enterprise	Enterprise	Backup	Backup	Enterprise	Enterprise
Encryption support	HW built-in	HW built-in	HW built-in	HW built-in	HW built-in	HW built-in
Buffer size	2 GB	2 GB or 4 GB	1 GB	1 GB or 2GB	2 GB	2 GB or 4 GB
Adaptive speeds	4	4	Dynamic	Dynamic	Dynamic 14	Dynamic
Price (typical street)	\$22,090	\$22,000 est	\$17,600 ⁴	\$17,600 est	\$20,000 est	\$20,000 est
Prev generations read	3	TBD	2	2	1	TBD
Prev generations written	0	0	1	1	1	TBD
Bit Error Rate (BER)	1x10 ⁻¹⁹	1x10 ⁻¹⁹	1x10 ⁻¹⁹	1x10 ⁻¹⁹	1x10 ⁻²⁰	1x10 ⁻²⁰
Drive manufacturers	1	1	3	3	1	1
Availability	2013	2017	2015	2017/2018	2014	TBD

² T10000D capacity can be increased to 8.5TB by setting a drive parameter, though actual capacity can vary, complicating tape-to-tape copies.

³ As reported by vendor or calculated by dividing tracks by channels.

⁴ Price cited is for a version compatible with the Oracle SL8500 robotic library. Standalone LTO7 drives cost approximately \$3,800.

Oracle T10000D

The T10000D is the Oracle flagship high-capacity enterprise tape drive typically used in conjunction with Oracle robotic libraries, such as the SL8500. EROS has eight T10000D drives installed in the SL8500 tape library.

Advantages

- The T10000D is an evolution of the T10000/T10000B/T10000C, which have performed reliably for the USGS.
- The Data Integrity Validation (DIV) feature provides internal block level checksum for server-less tape validation.
- Native capacity is 8.0 TB and native transfer rate is 252 MB/sec. By setting a parameter, 8.5 TB per tape may be possible. The T10000D can stream at multiple rates, which is important because some disks will not be able to keep up at 252 MB/sec. As tape speeds continue to increase, the disks must also.
- The T10000D is targeted to the enterprise storage market, where data viability, speed, and capacity are more important than cost.
- The T10000D was designed as a robust storage media, with the tape cartridge and drive built to withstand constant or frequent use in a robotic environment. The drives are compatible with the SL8500 and excel in a robotic environment because of their durability.
- T10000D drives provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- T10000D utilizes the same media as T10000C, written at a higher density – avoiding the expense of media replacement, which is a greater expense than drive replacement. The T10000D reads media written on the T10000/T10000B/T10000C/T10000D.

- The T10000D has a 2 GB buffer, which prevents occasional data starvation from reducing the transfer rate.
- The Bit Error Rate (BER) is 1×10^{-19} .
- Hardware encryption is built-in.
- Internal CRC ensures no data corruption on transfer.
- Drive partitioning allows positioning of data on a tape to improve access to critical data.
- Reclaim acceleration allows expired data to be overwritten.
- Supports Linear Tape File System (LTFS).

Disadvantages

- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs 144 passes (4,608 tracks divided by 32 channels). Multiple passes should not be a concern for archive operations because usage is limited.
- The T10000D drives cost 25 percent more than LTO7 drives and 10 percent more than the TS1150.
- T10000D sales are anticipated to be primarily for use in Oracle robotics. For this reason, the T10000D is anticipated to have a market share that will remain low compared to LTO, ensuring that media costs will remain high.
- The T10000D drive is only available from Oracle. This limited availability keeps the price high but does eliminate concerns of incompatibility.

Summary

The T10000D is a high-capacity, high-transfer rate, enterprise-class drive for use in Oracle robotic libraries. The cost of media and drives exceeds the cost of LTO, but media reuse for future

generations, shown in Figure 1, would effectively reduce media costs. The robust technology would be a prime choice for offline archives if only one copy of a dataset could be kept. When two or more copies of a dataset exist, and one is already on an enterprise technology, such as T10000C or T10000D, use of an enterprise solution for the second copy is not warranted.

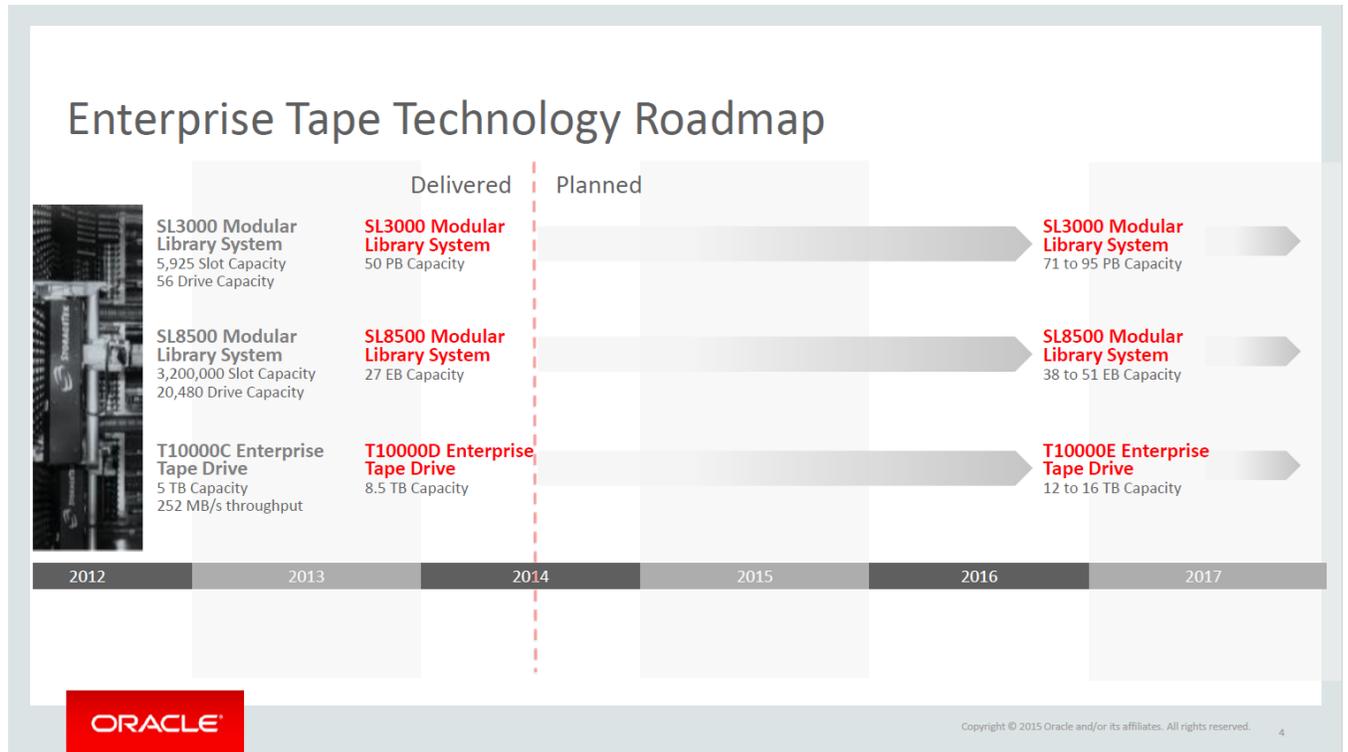


Figure 1. Oracle Roadmap (uncompressed) (Oracle, 2015).

Oracle T10000E

The T10000E is the fifth generation of the T10000 line. The T10000E is anticipated to ship in 2017.

Advantages

- The T10000E is an evolution of the T10000/T10000B/T10000C/T10000D, which have performed reliably for the USGS.
- The T10000E should have the Data Integrity Validation (DIV) feature which provides internal block level checksum for server-less tape validation.
- Native capacity is expected to be 12–16 TB and native transfer rate is 400–600 MB/sec. The T10000E is expected to stream at lower rates, which is important because most disks will not be able to keep up at 400+ MB/sec.
- The T10000E is expected to use at least 32 channels per pass. It is possible that Oracle will raise the number of channels to 64 to increase the transfer rate and reduce serpentine passes.
- The T10000E is targeted to the enterprise storage market, where data viability, speed, and capacity are more important than cost.
- The media for the T10000E is expected to be new, and the T10000E is expected to be able to read media written on the T10000B/T10000C/T10000D, and possibly the T10000.
- It is expected that the new T10000E drive and media will be designed to be robust, and built to withstand constant or frequent use in a robotic environment. The T10000E drives and media are expected to be compatible with the Oracle SL8500 robotic tape library.
- As with predecessor drives, T10000E drives should provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.

- The follow-on T10000F drives may use the same media, which would allow significant savings through media reuse.
- The T10000E is expected to have at least a 2 GB buffer, which prevents occasional data starvation from reducing the transfer rate.
- The BER is expected to be 1×10^{-19} or better.
- Hardware encryption is expected to be built-in.
- Expected to support Linear Tape File System (LTFS).

Disadvantages

- The T10000E drives are expected to cost more than the LTO and the TS1150.
- The T10000E is anticipated to be primarily for use in Oracle robotics. For this reason, the market share is anticipated to remain low compared to LTO.
- The T10000E drive is expected to be available only from Oracle. This availability keeps the price high but does eliminate concerns of incompatibility.
- The T10000E is expected to require new media in order to achieve the higher capacity, though it is also expected to be able to write T10000C/D media at T10000D density so that some older media could be retained by sacrificing capacity.
- Going forward, Oracle may not sell enough T10000 drives to recoup the cost of keeping up with IBM on head and media development. This may eventually lead to Oracle lagging IBM in capacity and performance.

Summary

Enterprise-class robustness of an offline archive copy is not required when the working copy of a dataset is already on enterprise-class technology in the EROS robotic library. The T10000E is not yet available and was therefore not assessed in the final evaluation.

IBM LTO7

The LTO7 is the most recent available generation of the LTO tape family. EROS has four LTO7 drives installed in the SL8500, in addition to several LTO5 and LTO6 drives. Several non-robotic LTO7 drives are also in use at EROS.

Advantages

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of Q4 2014, LTO held a 96 percent market share (Santa Clara Consulting Group, 2015). Several competing technologies, such as DLT (Betts, 2007) and SAIT, have been driven from the market.
- Native capacity is 6.0 TB and native transfer rate is 300 MB/sec.
- The LTO7 drive can adapt the transfer rate to match the streaming speed of the source.
- LTO7 is backward read compatible with LTO5 and LTO6, and backward write compatible with LTO6 (at the lower LTO6 density).
- LTO was developed by a consortium of HP, IBM, and Quantum (acquired from Seagate/Certance) and is licensed to others, including media manufacturers. This wide acceptance has introduced competition, which has, in turn, controlled costs. Compatibility tests are performed by the vendors.
- The LTO7 has a 1 GB buffer that prevents occasional data starvation from reducing the transfer rate.
- Hardware encryption is built-in.
- Supports Linear Tape File System (LTFS).

Disadvantages

- LTO is targeted to the backup market where speed, capacity, and cost are more important than long-term integrity of the data. Because backup tapes are write many/read rarely, errors would likely show up in a write pass where the errors can be worked around (rewrites) or the media discarded.
- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs 112 passes (3,584 tracks divided by 32 channels). Multiple passes should not be a concern for archive operations because usage is limited.
- Each generation of LTO requires new media to attain the rated capacity, ensuring that media costs will be higher until market saturation drives the price down. The price should not be a concern for archive operations, because required media life typically is supported by drive backward compatibility.
- LTO was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant enterprise/robotic use.
- IBM, HP, and Quantum (acquired from Seagate/Certance) co-developed LTO. This kind of partnership makes it possible for each vendor to interpret the specifications differently and to design drives that may have incompatibilities, though compatibility tests are performed.
- EROS experience with earlier LTO generations showed that drives slowly degrade prior to complete failure, resulting in slower transfer rates and marginal read capabilities. Substantial labor is required to monitor drives, perform problem analysis, re-archive data, and work with the vendor on drive replacement.

Summary

Staff completed tests of LTO7 technology at EROS in June 2016 (Gacke, 2016).

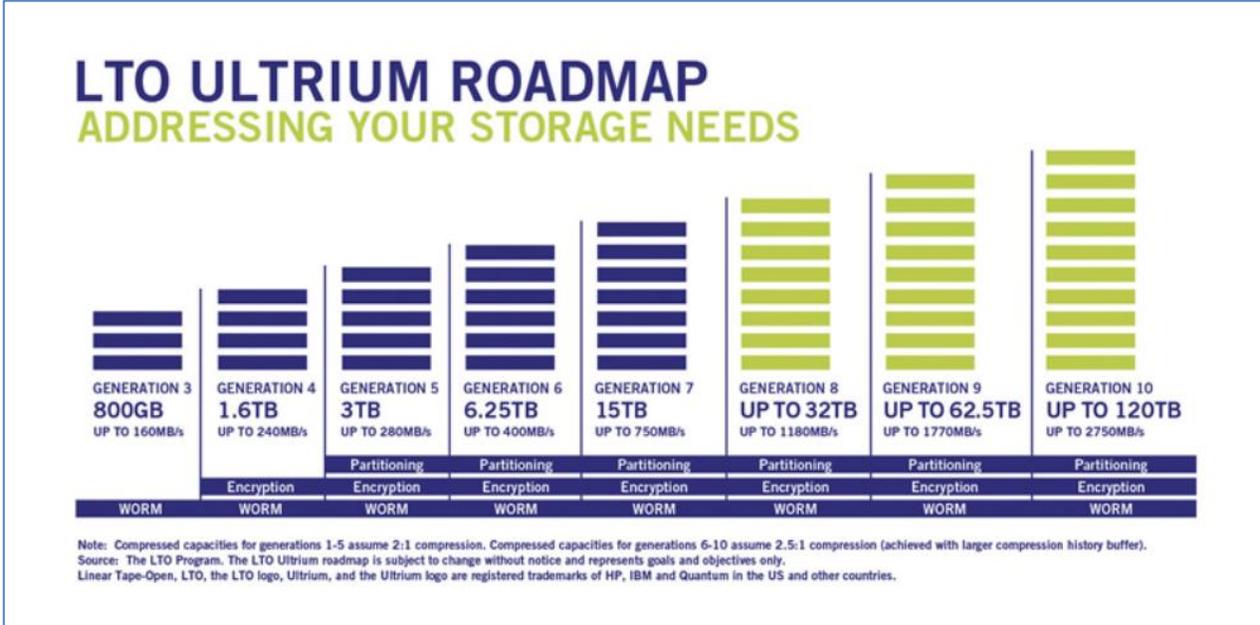


Figure 2. LTO Roadmap (compressed) (source: LTO Consortium).

IBM LTO8

The LTO8 is the next anticipated generation of the LTO tape family, with release expected in late 2017, or early 2018, based on a typical LTO release cycle of 2 years.

Advantages

- LTO has experienced phenomenal growth.
- Native capacity is expected to be 12.8 TB and native transfer rate is expected to be 472 MB/sec.
- The LTO8 drive is anticipated to use an adaptive transfer rate to match the streaming speed of the source.
- LTO8 should be backward read compatible with LTO6 and LTO7, and backward write compatible with LTO7 (at the lower LTO7 capacity).
- A consortium of HP, IBM, and Quantum (acquired from Seagate Certance) developed LTO; it is licensed to others, including media manufacturers. This wide acceptance introduced competition, which has, in turn, controlled costs. Compatibility tests are performed by the vendors.
- Hardware encryption is anticipated.
- Supports Linear Tape File System (LTFS).

Disadvantages

- LTO is targeted to the backup market, where speed, capacity, and cost are more important than long-term integrity of the data. Because backup tapes are write many/read rarely, errors would likely show up in a write pass where the errors can be worked around (rewrites) or the media discarded.

- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs at least 112 passes. Multiple passes should not be a concern for archive operations because use is limited.
- Each generation of LTO requires new media to attain the rated capacity, ensuring that media costs will be higher until market saturation drives the price down. The price should not be a concern for archive operations because required media life is typically supported by drive backward compatibility.
- LTO is a moderate use storage media, with the tape cartridge and drive not built to withstand constant use.
- IBM, HP, and Quantum (acquired from Seagate/Certance) co-developed LTO. This kind of partnership makes it possible for each vendor to interpret the specifications differently and to design drives that may have incompatibilities, though compatibility tests are performed.

Summary

Based on previous release intervals, LTO8 is expected to be announced in 2017 and made available in 2017 or 2018. LTO8 is not yet available and was not assessed in the final evaluation.

IBM TS1150

The TS1150 is an enterprise-class tape drive, used primarily in IBM robotic libraries and autoloaders. The TS1150 is a follow-on drive to the TS1140.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, TS1130, and TS1140.
- Supports dual 8-gigabit per second (Gbit/sec) Fiber Channel interfaces.
- Native capacity is 10 TB and native transfer rate is 360 MB/sec.
- The TS1150 is a robust storage technology, with the tape cartridge and drive built to withstand constant, or frequent, use in a robotic environment.
- The TS1150 uses some of the same media as earlier generations, plus a new higher-capacity cartridge.
- Hardware encryption is built-in.
- Supports Linear Tape File System (LTFS).
- The TS1150 is supported by IBM and Spectralogic tape libraries.

Disadvantages

- Designed primarily for use in IBM robotic libraries.
- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs 160 passes (5,120 tracks divided by 32 channels). Multiple passes should not be a concern for archive operations because use is limited.
- Not supported by the Oracle SL8500 robotic tape library currently in use at EROS.
- Less backward read capability than previous generations.

Summary

The TS1150 would not compare favorably in cost to LTO, and enterprise-class robustness is not required when the working copy of a dataset is already on enterprise-class technology in the EROS robotic library. In 2012, IBM announced it was preparing to demonstrate a 125 TB tape (Mellor, 2012).

	2005		2007		2009		2012	
	2003	2006	2008	2011	2014			
TS1100 Generations	Gen-1 3592 J1A	Gen-2 TS1120	Gen-3 TS1130	Gen-4 TS1140	Gen-5 TS1150	Gen-6	Gen-7	
Max Native Capacity	300 GB (JA)	700 GB (JB)	1.0 TB (JB)	4.0 TB (JC)	10 TB (JD)	15-20 TB (JD)	30-40 TB	
Other Native Capacities with Media Reuse		500 GB JA	640 GB JA	1.6 TB JB 640 GB JA R/O	7 TB JC	8-10 TB JC 15-20 TB JD	8-10 TB JC 15-20 TB JD 30-40 TB JE	
Native Data Rate	40 MB/s	100 MB/s	160 MB/s	250 MB/s	Up to 360 MB/s	Up to 540 MB/s	1000 MB/s	

Figure 3. IBM Roadmap (uncompressed) (Source: IBM)

IBM TS1160

The TS1160 is anticipated to be the next generation of the 3592 tape family, with release expected in 2017 based on frequency of past releases. Note that the TS1160 name has not been confirmed, but follows logically.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, TS1130, TS1140, and TS1150.
- Should support at least two 8-Gbit/sec Fiber Channel interfaces.
- Native capacity is expected to be 20 TB and native transfer rate may reach 720 MB/sec.
- The TS1160 will be a robust storage technology, with the tape cartridge and drive built to withstand constant, or frequent, use in a robotic environment.
- The TS1160 may use some of the same media as the TS1150.
- A hardware encryption feature should be included in the drive.
- Expected to support Linear Tape File System (LTFS).

Disadvantages

- Designed primarily for use in IBM robotic libraries.
- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs at least 160 passes. Multiple passes should not be a concern for archive operations because usage is limited.

Summary

The TS1160 would not compare favorably in cost to LTO, and enterprise-class robustness is not required when the working copy of a dataset is already on enterprise-class technology in the EROS

robotic library. In 2012, IBM announced preparations to demonstrate a 125 TB tape (Mellor, 2012). TS1160 is not yet available and was not assessed in the final evaluation.

Tables

3.1 Design criteria

The design criteria and target market of a drive are interrelated. LTO7 is targeted to the backup market, as demonstrated by LTO marketing. The T10000D and TS1150 are targeted to the enterprise (data center) market.

A drive targeted to the backup market is designed for write many/read rarely and depends on write error detection because the data are still available and can easily be rewritten. Backup drives are typically built for speed, capacity, and low cost.

A drive targeted to the enterprise market is designed for write many/read many use in a robotic library, and equal emphasis is placed on detecting errors on read and write. Enterprise drives are typically built for reliability and speed, with capacity as a secondary factor. Cost is a not a primary consideration to enterprise users willing to pay for quality.

A drive targeted to the archival market would be designed for write once/read rarely, and more emphasis would be placed on detecting and correcting errors on read; however, there are currently no drives designed or marketed primarily for archive use.

The following formula ranked design criteria:

$$\begin{aligned} & ((200 - \text{serpentine passes}) / 10) + \\ & (\text{absolute value of error rate exponent} / 2) + \\ & (\text{construction } 3 = \text{moderate usage, } 5 = \text{high usage}) + \\ & (\text{head contact } 3 = \text{contact, } 5 = \text{min contact}) \end{aligned}$$

/ 2.51 (to adjust the highest rank to 10)

Table 4. Design Criteria and Target Market.

Technology	Serpentine Tracks/ Passes	Target Market	Tape Composition	Uncorrected Error Rate	Cartridge Construction Rating	Head Contact	Ranking
Oracle T10000D	4608/144	Enterprise	BaFe	1×10^{-19}	High usage	Minimum contact	10.0
IBM LTO7	3584/112	Backup	MP or BaFe	1×10^{-19}	Moderate usage	Contact	9.7
IBM TS1150	5120/160	Enterprise	BaFe	1×10^{-20}	High usage	Contact	8.8

3.2 Transfer Rate

Transfer rate is important because it establishes how quickly the migration and verification of an archive dataset may be completed and how fast a recovery can be completed. The minimum read transfer rate requirement is 200 MB/sec, with 250 MB/sec desired. Much of the data archived at the USGS are raster imagery that typically lacks repeatable patterns that would compress well; therefore, all transfer rates cited are native (uncompressed).

Where measured transfer rates were not available, approximate rates are determined based on the accuracy of specified transfer rates of previous generations. The source of the transfer rates are noted in table 5.

The ranking was determined by adding the actual/approximate read and write rates for each drive, setting the ranking for the fastest drive to 10, then ranking the others against the leader. For example, a drive having one-half of the total read/write transfer rate of the leader would be ranked 5.

Table 5. Transfer Rates.

Tape Drive Technology	Advertised Native Rate	Source of Test Results	Actual/approximate Native Write Transfer Rate	% of Advertised Rate	Actual/approximate Native Read Transfer Rate	% of Advertised Rate	Ranking
Oracle T10000D	238 MB/sec	EROS testing	241 MB/sec	95.2%	241 MB/sec	95.6%	6.7
IBM LTO7	300 MB/sec	EROS testing	271 MB/sec	90.3%	274 MB/sec	91.3%	7.6
IBM TS1150	360 MB/sec	Vendor	360 MB/sec	100%	357 MB/sec	99.2%	10.0

3.3 Capacity

A secondary requirement is to conserve rack or pallet storage space and reduce tape handling by increasing per-media capacity. The current archive media of choice at the USGS is LTO7 at 5.7 TB of usable capacity per tape. The minimum capacity requirement is 4 TB, with 5 TB or more desired. All three reviewed technologies meet the 4 TB requirement based on the advertised capacity. Because much of the data archived are not compressible, all capacities are native (uncompressed). Where measured capacities were not available, approximate capacities were determined based on the accuracy of specified capacities of previous generations.

The capacities listed in Table 6 presume that a gigabyte equals 1,073,741,824 bytes. The ratings were determined by computing each actual, or approximate, capacity score as a percentage of the highest capacity drive on a scale of 1 to 10, with the highest capacity as a 10. The source of the capacity ratings are noted in Table 6. Capacity yield varies by media vendor.

Table 6. Storage Capacities.

Tape Drive Technology	Advertised Native Capacity	Actual/approximate Native Capacity	% of Advertised Capacity	Ranking
Oracle T10000D	8.5 TB	8.13 TB EROS measured	95.7%	8.6
IBM LTO7	6.0 TB	5.70 TB EROS measured	95.0%	6.0
IBM TS1150	10.0 TB	9.50 TB approximate	95.0% approximate	10.0

3.4 Cost Analysis

Table 7 shows the relative drive and media costs, and the cost per terabyte for media. Rankings were established by setting the least expensive (drive and media) to 10, then rating each of the others against the lowest cost. Media costs per terabyte are based on advertised capacity. Costs do not include system interfaces or cables. Prices are based on the lowest price present on the Web or on government price lists or recent actual discount prices, and include one year warranty or support.

Unlike LTO, Oracle and IBM drives have historically allowed media to be written across two or more generations of drives, increasing tape capacity on the newer drives. This advantage is not depicted in the following table because writing archive tapes is usually a one-time operation before archive tapes are shipped offsite permanently, which would not allow for taking advantage of higher capacity with newer drives. The capability to rewrite media at higher capacity would make a case for these technologies to be used for nearline or onsite offline copies, because the media would be readily accessible.

Table 7. Drive and Media Costs

Tape Drive Technology	Drive \$/unit	Media \$/unit	Media \$/TB	Ranking Drive Cost	Ranking Media Cost/TB
Oracle T10000D	\$22,090	\$218	\$26	7.8	7.7
HP LTO7	\$17,600	\$118	\$20	10.0	10.0
IBM TS1150	\$20,000	\$313	\$31	8.8	6.5

3.5 Scenarios

The total drive and media cost for three scenarios is shown in Table 8. These scenarios presume that each dataset or project stands alone, although pooling resources for multiple datasets can mitigate cost. Where there is market competition, a considerable drop in media prices often occurs within 6 months after product introduction.

Rankings are based on the 400 TB option and were established by setting the least expensive to 10, and then rating each of the others against the lowest cost. Advertised native capacities are used. Costs do not include system interfaces or cables, but do include one year of maintenance on the drives.

Though not represented in this study, technology refresh costs related to moving from one generation to the next may vary depending on whether the vendor requires a media change. LTO has always required new media for each generation, but Oracle and IBM have historically used the same media for at least two generations.

Table 8. Scenario Costs (drives, media)

Technology	200 TB 2 Drives	400 TB 3 Drives	1000 TB 4 Drives	400 TB Ranking
Oracle T10000D	\$49,512	\$76,734	\$114,520	7.9
IBM LTO7	\$39,212	\$60,824	\$90,460	10.0
IBM TS1150	\$46,260	\$72,520	\$111,300	8.4

3.6 Vendor Analyses

An analysis of each company and the stability of each technology is shown in Table 9. All are established and stable companies; therefore, this rating should not be viewed as a market analysis. When selecting an archive technology, it makes sense to look at the company and product histories even though rating vendor history is challenging because of mergers and acquisitions. For T10000D, the technology was based on the ancestor 9940; therefore, the technology age includes the 9940. The longevity rankings were determined by the following formula:

$$(\text{company age} + \text{technology age}) / 12.6 \text{ (to adjust the highest rank to 10)}$$

Determining company years in business is complicated by mergers and acquisitions, such as when Sun acquired STK and was later acquired by Oracle. The years in business began with STK because the tape technology offered today is based on STK products. The purpose of this section is to assess technology lineage and company history, but mergers and acquisitions may be distractive and detrimental when considering lineage and history.

Table 9. Vendor Analyses

Company	Technology	Years in Business	Technology Age in Years	Longevity Ranking
Oracle/Sun/STK	T10000	47 (1969)	16 (2000)	5.0
IBM	LTO	105 (1911)	16 (2000)	9.6
IBM	3592 (3590)	105 (1911)	21 (1995)	10.0

3.7 Drive Compatibility

The level of intergeneration drive compatibility and planned future drives are shown in Table 10. The column “Previous Generations Read” indicates backward read compatibility. Backward write compatibility is of little consequence for archiving, so is not considered. The column "Future Generations Mapped" indicates the number of generations planned in the current drive family, following the drive indicated. The following formula determined the ranking:

$$(\text{Previous Generations Read} + \text{Future Generations Planned}) \times 2 \text{ (to adjust the highest rank to 10)}$$

Table 10. Drive Compatibility

Technology	Previous Generations Read	Future Generations Mapped	Ranking
Oracle T10000D	3	1	8.0
IBM LTO7	2	3	10.0
IBM TS1150	2	2	8.0

3.8 Ranking Summary

The ranking summary provides a quick reference to the rankings.

Table 11. Ranking Summary (blue indicates the highest ranking in category)

Drive	Design Criteria	Capacity	Media Cost	Drive Compatibility	Transfer Rate	Drive Cost	Vendor Analyses	Scenario Cost
T10000D	10.0	8.6	7.7	8.0	6.7	7.8	5.0	7.9
IBM LTO7	9.7	6.0	10.0	10.0	7.6	10.0	9.6	10.0
IBM TS1150	8.8	10.0	6.5	8.0	10.0	8.8	10.0	8.4

Conclusions and Recommendations for USGS Offline Archiving Requirements

4.1 Weighted Decision Matrix

A weighted analysis of the drives considered is shown in Table 12. The criteria emphasize the importance of traits contributing to data preservation. The USGS made the final decision regarding which criteria to use and the relative weighting of the criteria. The columns with a green heading are relative ratings for each technology. The columns with a yellow heading are calculated by multiplying the relative weight by the relative rating. The following list describes each criterion:

- Design (reliability of media) — this criterion describes the ability of the media to remain readable with time. Included in this criterion is the number of passes per full-tape read or write, cartridge construction, uncorrected BER, and amount of head contact (Table 4).
- Capacity — this criterion describes the measured or approximate capacity per cartridge, which is typically less than the advertised capacity (Table 6).
- Media Cost/TB — this criterion is a rating of the relative cost per terabyte for media using the advertised capacity (Table 7).
- Compatibility — this criterion describes the likelihood that the drive technology will continue to evolve and the extent to which future drives will have backward read capability. This criterion will give an indication of the ability to maintain drives that can read an aging archive (Table 10).
- Transfer Rate — this criterion describes the aggregate read and write transfer rate, which is typically less than the advertised transfer rate (Table 5).
- Drive Cost — this criterion is the rating of relative cost of each drive at the lowest currently available price (Table 7).

- Vendor Analyses — this criterion is the rating of the viability of the vendor and technology (Table 9).
- Scenario cost — this criterion is the rating of the cost of scenario #1, which comprises media cost and drive cost. The advertised capacity is used (Table 8).

In the decision matrix spreadsheet shown in Table 12, not all criteria have been selected for the final analysis of this trade study. These unused criteria were provided in the spreadsheet so that users may insert the criteria weights for their specific application.

Table 12. Weighted Decision Matrix

Selecton Criteria	Wt	Oracle T10000D	IBM LTO7	IBM TS1150	Oracle T10000D	IBM LTO7	IBM TS1150
Design criteria		10.0	9.7	8.8	0.0	0.0	0.0
Capacity	20	8.6	6.0	10.0	172.0	120.0	200.0
Media cost/TB		7.7	10.0	6.5	0.0	0.0	0.0
Compatibility	15	8.0	10.0	8.0	120.0	150.0	120.0
Transfer rate	15	6.7	7.6	10.0	100.5	114.0	150.0
Drive cost		7.8	10.0	8.8	0.0	0.0	0.0
Vendor analyses	15	5.0	9.6	10.0	75.0	144.0	150.0
Scenario cost	35	7.9	10.0	8.4	276.5	350.0	294.0
Total Weighted Score					744.0	878.0	914.0

4.2 Conclusions and Notes

- TS1150 achieved the highest total score in the study; however, no compelling reason exists to abandon LTO to adopt a new standard offline archive technology based solely on these relatively close scores. The TS1150 would be incompatible with the existing Oracle SL8500 robotic library.
- A TS1150 drive was not available to be tested for this study; therefore, performance and capacity figures were based on vendor benchmarks where available or on drive specifications combined with past performance (percentage of the claimed specifications that were achievable in the past).
- When multiple copies of a dataset are maintained, trading cost and performance for reliability is acceptable, particularly when the working copy is on an enterprise technology, such as Oracle T10000D, as are most archives at EROS.
- Using multiple tape technologies, where multiple copies of an archive exist, could mitigate the risk that one technology has an unforeseen drive or media issue.
- As any drive saturates the market, media and drive costs drop. Based on EROS experience with enterprise tape technology and observation of Oracle and IBM pricing, enterprise drives such as the T10000D and TS1150 are unlikely to achieve a level of market saturation that would cause substantial price decreases.
- With proper handling and multiple copies, any of the technologies evaluated in this report could be deployed for archive use. When more than two copies exist, all could be on non-enterprise technology.

4.3 Recommendations

Although the IBM TS1150 scored highest in this study, there is no compelling reason to adopt a new standard archive device at this time. The TS1150 would not be compatible with the existing Oracle SL8500 robotic tape library. Automation would require the purchase of an IBM tape library, and substantial engineering effort would be required to integrate the two storage systems.

1. The USGS should continue with LTO7 as the offline storage media of choice, then test and move to LTO8 when available.
2. Data archived on LTO5 should be migrated to LTO7.
3. To reduce risk, the USGS should continue the strategy of storing datasets on multiple technologies. For example, store a working copy of a dataset on nearline T10000D, store offline/onsite data on LTO7 or T10000D, and store offline/offsite copies on LTO7. This strategy partly mitigates the risks of one or the other technology failing or being retired prematurely.
4. In addition to a nearline and offsite copy of a dataset, an onsite offline copy should be maintained, providing fast recovery without risking the shipping of the offsite LTO copy. This practice has been implemented and should continue.
5. The USGS periodically tests archive tapes for readability, which should be continued. This testing should not be extensive enough to incur undue wear on the media or frustrate NARA, but should be frequent enough to provide an opportunity to detect deteriorating media. It is recommended to retrieve a box of media once per year from NARA so that a specific set of files can be tested. This practice has been implemented and should continue.

6. All archived files should be checksummed, with the checksum stored in the corresponding inventory record. When a file is retrieved from either the HSM or the offline media, integrity can be verified. Verification of each retrieved file may not be feasible because of CPU impacts. This practice has been substantially implemented and should continue. Vendor verification features could be utilized instead of or in addition to file checksums.
7. All data should be migrated to new media approximately 5 years after written. Although most tape technologies can reliably store data for much longer periods, after 5 years the transfer rates and densities that once were leading edge will become problematic, and drives will become difficult to maintain.
8. As archive media is retired it should be sent back to EROS from NARA and 10% of the media read before disposal, to assess storage conditions and media viability. This practice has been implemented and should continue.
9. When writing archive tapes, the tapes should be verified on a second drive. This verification will help identify any drive incompatibility. This practice has been implemented and should continue.
10. Where possible, the USGS should avoid buying media brands that have proven unreliable.
11. The USGS should plan to periodically update this trade study. Annual updates may be too frequent to observe market changes because drives are typically updated on a 2- or 3-year cycle. Each time this study is revisited, the highest scoring technology may change. This change does not indicate that the USGS should change offline tape technologies frequently. Staying with a given technology for several years is beneficial, even if the technology is not continuously the leading technology. This study is a snapshot in time, and results would differ even a few months earlier/later because of new hardware releases.

References Cited

- Betts, B. (2007, March 21). *DLT tape reaching end of reel*. Retrieved from The Register:
http://www.theregister.co.uk/2007/03/21/lto_beats_dlt/
- Gacke, K. (2016). *LTO TAPE DRIVE PERFORMANCE*.
- Mearian, L. (2010, January 22). *Fujifilm, IBM tech could yield 35TB magnetic tape drives*. Retrieved from computerworld.com:
http://www.computerworld.com/s/article/9147498/Fujifilm_IBM_tech_could_yield_35TB_magnetic_tape_drives
- Mellor, C. (2012, Oct 12). *IBM Prepares to demo 125TB MONSTER tape*. Retrieved from TheRegister:
http://www.theregister.co.uk/2012/10/19/ibm_125tb_tape/
- OSTA. (2003). *Understanding CD-R & CD-W*. Retrieved from http://www.osta.org/technology/pdf/cdr_cdrw.pdf
- Santa Clara Consulting Group. (2015). *Worldwide Tape Market Analysis for 4Q14*. Retrieved from storagenewsletter.com:
<http://www.backupworks.com/worldwide-tape-cartridge-market-LTO-Ultrium.aspx>
- SearchStorage. (n.d.). *HVD (holographic versatile disc)*. Retrieved from SearchStorage:
<http://searchstorage.techtarget.com/definition/HVD>
- Sony. (2014, April 30). *Sony develops magnetic tape with the world's highest areal recording density of 148 Gb/in²*. Retrieved from Sony.net: <http://www.sony.net/SonyInfo/News/Press/201404/14-044E/index.html>

Appendix: Supplemental Information

Vendor Sites

<http://www8.hp.com/us/en/products/tape-drives-enclosures/> (HP)

<http://www.oracle.com/us/products/servers-storage/storage/tape-storage/index.htm> (Oracle)

<http://www-03.ibm.com/servers/storage/tape/index.html> (IBM)

https://www.ibm.com/support/knowledgecenter/STCMML8/com.ibm.storage.ts3500.doc/ipg_3584_a69p0cti20.html (IBM)

https://www.ibm.com/support/knowledgecenter/STQRQ9/com.ibm.storage.ts4500.doc/ts4500_ipg_drives_3592.html (IBM)

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM+Removable+Media+on+IBM+i/page/3592+Cartridges+and+Drives+compatibility+on+IBM+i> (IBM)

https://www.ibm.com/support/knowledgecenter/STZMZN/com.ibm.storage.hollywood.doc/lfs_supported_tape_drives_library_edition.html (IBM)

<http://www.quantum.com/Products/TapeDrives/Index.aspx> (Quantum)

Consortium Sites

<http://www.lto.org/>

Other

<http://www.infoworld.com/article/3040522/storage/samsung-ships-the-worlds-highest-capacity-ssd-with-15tb-of-storage.html>

<http://www.slideshare.net/JosefWeingand/tsm-symposium-2015-ibm-future-of-tape>

<http://www.lascon.co.uk/Tape-drive-comparisons.php>

<http://www.slideshare.net/spectrallogic/ts1150-webinar-slides>

https://www2.cisl.ucar.edu/sites/default/files/Raymond_0.pdf

<https://www.engadget.com/2015/09/28/first-nano-photonic-memory-chip/>

<http://storageconference.us/2016/Slides/AprilAlstrin.pdf>

https://gcn.com/articles/2015/04/13/ibm-tape-storage.aspx?s=gcntech_140415

<https://www.tapeandmedia.com/lto-7-tape-media-tapes.asp>

http://www.infostor.com/index/blogs_new/dave_simpson_storage/blogs/infostor/dave_simpson_storage/post987_3968595795243568175.html

<http://www.drunkendata.com/?p=4330>

http://www.theregister.co.uk/Print/2013/09/25/the_state_of_tape/

<http://highscalability.com/blog/2011/3/24/strategy-disk-backup-for-speed-tape-backup-to-save-your-baco.html>

http://www.theregister.co.uk/2010/08/11/oracle_on_storage/

<http://www.engadget.com/2010/02/08/inphase-out-of-business-assets-seized-for-back-taxes/>

http://en.wikipedia.org/wiki/Holographic_Versatile_Disc

http://www.theregister.co.uk/2007/03/21/lto_beats_dlt/

<http://www-03.ibm.com/press/us/en/pressrelease/29245.wss>

http://en.wikipedia.org/wiki/Linear_Tape-Open

<http://www.infoworld.com/d/data-explosion/tape-dead-long-live-tape-090>