TECHNICAL SUPPORT SERVICES CONTRACT (TSSC)		
USGS EROS	G15PC00012	DELIVERABLE ID: 1.0007-1.110.1.W01.1.2.6-A4704
FY18 OFFLINE ARCHIVE MEDIA TRADE STUDY	DATE: 07/31/18	VERSION NBR: 1.0

ARCHIVE AND RECORDS MANAGEMENT

FISCAL YEAR 2018 ARCHIVE MEDIA TRADE STUDY

This delivery is official submission to the USGS under the TSSC contract G15PC00012.

TSSC FISCAL YEAR 2018 OFFLINE ARCHIVE MEDIA TRADE STUDY

Remote Sensing, CEOS, and Archives Coordination Project

By Tom Bodoh[1]

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, Randall Sunne, Brian Reutzel, and Mark Barber.

---

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

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 2018 revision of a study completed in Fiscal Year 2001 and revised in Fiscal Years 2003, 2004, 2006, 2008, 2010, 2012, 2014, and 2016.

Contents

Figures

Tables

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-Packard
HSM	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
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
TMR	Tunnel Magnetoresistance
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.
- Revised 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-year 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.

July 2018

- Overall refresh of study, removing most references to older technologies.
- Changed LTO8 and TS1155 to current technologies.

- Added LTO9 as a future technology.
- Oracle T10000D changed to non-evaluated status since Oracle has ceased further tape drive development. For this iteration of the study, the T10000D will be left in the specification table since it is still sold.
- Cloud was added to the specification table, but is a dismissed and non-evaluated technology in this iteration of the study.
- In the Vendor Analyses criteria calculation, adjusted the start of the IBM TS11 technology line to begin with the model 3592 tape drive in 2003 instead of the model 3590 tape drive in 1995.

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 chance 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.

The 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 system.

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, continues to archive 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 (currently used media 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	**2003–2007**	**200 GB**	**40.0 MB/sec**	**Digital**
HP LTO Ultrium 3	**2005–2010**	**400 GB**	**80.0 MB/sec**	**Digital**
HP LTO Ultrium 4	**2007–2015**	**800 GB**	**120.0 MB/sec**	**Digital**
HP LTO Ultrium 5	**2010–2017**	**1.5 TB**	**140.0 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
IBM LTO Ultrium 8	2018-present	12.0 TB	360.0 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 use 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 typically 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 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 its 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 which has ongoing development and a roadmap. 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 8 terabytes (TB) [8,000 gigabytes (GB)] of uncompressed data.
3. The technology must have an uncompressed write transfer rate of at least 300 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 were proven to be unreliable, and are no longer available.

The following currently available drive technologies were selected for consideration:

1. LTO8 (Linear Tape Open)— International Business Machine (IBM) representative of models by Quantum and Hewlett Packard Enterprise (HPE).
2. IBM TS1155

The following future drive technologies are mentioned but not considered since they are not available yet:

1. IBM LTO9
2. 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 costlier 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 since it can become unusable if stiction keeps a parked head adhered to the disk surface, preventing spin-up.

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 will continue to gradually replace magnetic disk. 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), Blu-Ray, 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 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. Blu-Ray does not meet the evaluation criteria for this study. The market for optical disc is driven by entertainment but is dwindling due to the entertainment industry transitioning to 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. In 2016, Sony released its petabyte-scale Everspan Optical Robotic Library. The website recently shut down, and there is no recent information on it.

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 and IBM of a magnetic substrate technology, which could result in a tape product with a capacity of up to 330 TB

(Sony, 2017). There is no indication if, or when, a tape drive will become available, or if it will be a Sony or IBM product.

1.5.5 Cloud Storage

Though not an offline media, cloud storage is emerging as a potential offsite storage alternative, which could be used as one copy of an archive. At the present time, online public-cloud storage has significant cost considerations for petabyte-scale datasets, but could be leveraged as a working copy of limited datasets. Other considerations include data integrity, location, security, and contract termination. Public-cloud storage would currently cost considerably 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.

Technical Assessment

2.1 Analysis

This technical assessment includes drives selected for final evaluation (LTO8 and TS1155) and drives anticipated to be released in the next two years (LTO9 and TS1160). LTO drives are available from multiple vendors (Quantum, IBM, and HP), with an IBM drive selected to represent LTO technology in this study. The following offline storage technologies will be evaluated, but only the drives shown in bold will be included in the analysis and final evaluation:

- Oracle T10000D
- **IBM LTO8**
- IBM LTO9
- **IBM TS1155**
- IBM TS1160
- AWS Glacier

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

Specification	T10000D	LTO8	LTO9	TS1155	TS1160	AWS Glacier
Uncompressed capacity	8.5 TB[2]	12 TB	24 TB	15 TB	20 TB	Unlimited
Uncompressed xfer rate	252 MB/sec	360 MB/sec	<450 MB/sec	360 MB/sec	<500 MB/sec	Unknown
Recording technology	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine	Unknown
Tracks	4,608	6,656		7,680		NA
Channels	32	32	32 or 64	32	32 or 64	NA
Passes[3]	144	208		240		NA
Tape velocity (read)	4.75 m/sec	4.731 m/sec		6.22 m/sec		NA
Type	Enterprise	Backup	Backup	Enterprise	Enterprise	Enterprise
Encryption support	HW built-in	HW built-in	HW built-in	HW built-in	HW built-in	Automatic
Buffer size	2 GB	1 GB	1 GB or 2GB	2 GB	2 GB or 4 GB	NA
Adaptive speeds	4	Dynamic	Dynamic	Dynamic 12	Dynamic 12	NA
Price (typical street)	$22,090	$19,200[4]	$19,200 est	$20,000 est	$20,000 est	$.004/GB/Mo+
Prev generations read	3	1	2	1	TBD	NA
Prev generations written	0	1	1	1	TBD	NA
Bit Error Rate (BER)	$1x10^{-19}$	$1x10^{-19}$	$1x10^{-19}$	$1x10^{-20}$	$1x10^{-20}$	Unknown
Drive manufacturers	1	3	3	1	1	NA
First availability	2013	2017	2019/2020	2017	2019/2020	2012

[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] Competing brands may be considerably lower cost.

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 twelve T10000D drives installed in the SL8500 tape library. Oracle has discontinued further development of tape-drive technology, specifically the T10000 line. It is anticipated that in the near future, the T10000D tape drive will be discontinued, though likely supported for 5 additional years. After T10000D end-of-support, LTO will be the only supported tape technology for use in Oracle tape libraries, such as the SL8500.

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, disk speed must keep up.
- 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 $1x10^{-19}$.
- 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

- In early 2017, Oracle halted further development of tape drives, including the T10000 line, and it is anticipated that sales of T10000D drives will soon be discontinued. End of support is typically at least five years after end of availability.
- 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 use is limited.
- The T10000D drives cost 15 percent more than LTO8 drives and 10 percent more than the TS1155.

- T10000D sales are 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 potential use for an offline archive media is negated by the lack of future development of the T10000 line, as evidenced by the LTO8 quickly eclipsing the stagnant T10000D in capacity, transfer rate, and price, and matching the BER.

IBM LTO8

The LTO8 is the most recent available generation of the LTO tape family. EROS has four LTO8 drives on order for the SL8500, in addition to several existing LTO5, LTO6, and LTO7 drives. Several non-robotic LTO8 drives are planned for FY19.

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, and it appears that T10000 is next.
- Native capacity is 12.0 TB and native transfer rate is 360 MB/sec.
- The LTO8 drive can adapt the transfer rate to match the streaming speed of the source.
- LTO8 is backward read compatible with LTO7, and backward write compatible with LTO7 (at the lower LTO7 density).
- LTO was developed by a consortium of HP, IBM, and Quantum 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 LTO8 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).
- The BaFe tape media formulation used by LTO7 and later drives has proven much more stable than the previous MP media, minimizing labor for problem analysis.

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 208 passes (6,656 tracks divided by 32 channels). 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.
- LTO was designed as a moderate use storage media, with the tape cartridge and drive not built to withstand constant enterprise/robotic use.
- IBM, HP, and Quantum co-develop 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.
- Due to a new head technology designed to support the higher densities allowed by BaFe tape formulation and incompatible with the MP tape formulation of LTO6, the LTO8 is backward read compatible with only one generation (LTO7), not two generations as with LTO3

through LTO7. Indications are that LTO9 will again be backward read compatible with two generations (LTO7 and LTO8).

Summary

LTO8 drives will not be available for testing at EROS until after this study will be completed, so estimated capacity and transfer rates will be calculated based on LTO7 actual vs. specification percentage.

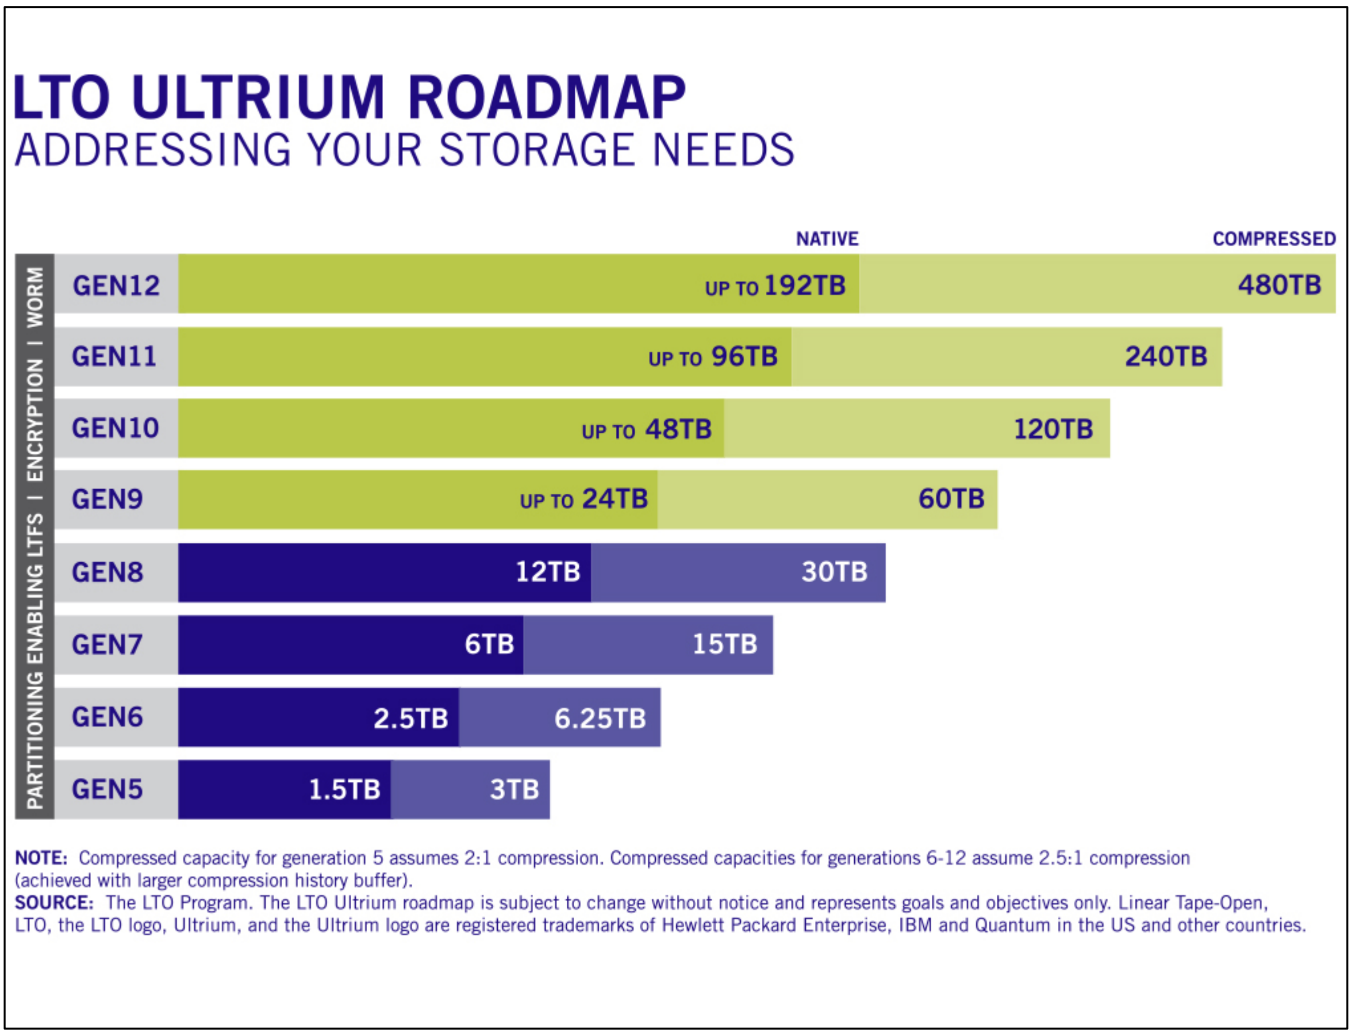

Figure 1. LTO Roadmap (source: LTO Consortium).

IBM LTO9

The LTO9 is the next anticipated generation of the LTO tape family, with release expected in late 2019, or early 2020, based on a typical LTO release cycle of 2 years.

Advantages

- LTO has experienced phenomenal growth.
- Native capacity is expected to be 24 TB and native transfer rate is expected to be 400 to 450 MB/sec.
- The LTO9 drive is anticipated to use an adaptive transfer rate to match the streaming speed of the source.
- LTO9 should be backward read compatible with LTO7 and LTO8, and backward write compatible with LTO8 (at the lower LTO8 capacity).
- LTO was developed by a consortium of HP, IBM, and Quantum 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.
- Hardware encryption is anticipated.
- Linear Tape File System (LTFS) support is anticipated.
- The BaFe tape media formulation utilized by LTO7 and later drives has proven much more stable than the previous MP media, minimizing labor for problem analysis.

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 208 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.
- LTO is a moderate use storage media, with the tape cartridge and drive not built to withstand constant use.
- IBM, HP, and Quantum co-develop 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, LTO9 is expected to be announced in 2019 and made available in 2019 or 2020. LTO9 is not yet available and was not assessed in the final evaluation.

IBM TS1155

The TS1155 is an enterprise-class tape drive, used primarily in IBM robotic libraries and autoloaders. The TS1155 is a follow-on drive to the TS1150.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, TS1130, TS1140, and TS1150.
- Supports dual 8-gigabit per second (Gbit/sec) Fiber Channel interfaces or dual 10 Gbit/sec Ethernet interfaces.
- Native capacity is 15 TB and native transfer rate is 360 MB/sec.
- The TS1155 is a robust storage technology, with the tape cartridge and drive built to withstand constant, or frequent, use in a robotic environment.
- The TS1155 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 TS1155 is supported by IBM and Spectra Logic tape libraries.

Disadvantages

- Compatible with IBM and Spectra Logic robotic libraries, but not currently supported in Oracle tape libraries such as the SL8500 in use at EROS.
- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs 240 passes (7,680 tracks divided by 32 channels). Multiple passes should not be a concern for archive operations because use is limited.

Summary

The IBM TS1155 drive compares favorably in cost to the IBM LTO8 drive, though enterprise-class robustness is not required when the working copy of a dataset is already on enterprise-class T10000 technology in the EROS robotic library. In 2017, IBM and Sony announced they had co-created a 330 TB tape (Seppala, 2017).

IBM TS1160

The TS1160 is anticipated to be the next generation of the 3592-tape family, with release expected in 2019 or 2020, based on frequency of past releases.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, TS1130, TS1140, TS1150 and TS1155.
- Should support dual 16-gigabit per second (Gbit/sec) Fiber Channel interfaces or dual 25 Gbit/sec Ethernet interfaces.
- Native capacity is expected to be 20 TB and native transfer rate between 420 and 500 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 and TS1155.
- A hardware encryption feature should be included in the drive.
- Expected to support Linear Tape File System (LTFS).

Disadvantages

- Compatible with IBM and Spectra Logic robotic libraries, but not currently supported in Oracle tape libraries, such as the SL8500 in use at EROS.
- Frequent end-to-end use of a tape would be a concern because one end-to-end read/write incurs at least 240 passes. Multiple passes should not be a concern for archive operations because use is limited.

Summary

The IBM TS1160 is expected to compare favorably in cost to the IBM LTO9 drive, though enterprise-class robustness is not required when the working copy of a dataset is already on enterprise-class T10000 technology in the EROS robotic library. TS1160 is not yet available and was not assessed in the final evaluation.

AWS Glacier

Amazon Web Services (AWS) is the leading cloud vendor by far. In 2012, AWS released Glacier as its deep-archive service. Comparing a storage service to storage media is difficult, and the results of the comparison may not be useful.

Advantages

- The service has been in operation since 2012 and there is no indication that it will be retired in the near future.
- Capacity is theoretically unlimited as AWS has constantly grown total capacity, apparently staying ahead of demand. A customer with exabytes of requirement would almost certainly need to wait for resources to be added, but a planned move of petabytes of data should be possible if plans are shared with AWS in advance.
- Glacier and competing archive storage could provide a strategically-placed copy of very infrequently used data, if the application is also on the cloud.
- All data on Glacier is encrypted.
- All data on Glacier is stored with redundancy.
- Migration to new media is automatic, at no charge, and disposal of the old media is on AWS.

Disadvantages

- Data recovery from a public-cloud archive, such as Glacier, would present challenges, such as retrieval fees, egress fees, and potentially long delays.
- AWS glacier storage fees of $.004/GB/mo are lower than some competitors, but still substantial for large archives. An archive containing 20 PB (petabytes) would cost $960,000/yr, not including retrieval fees or egress fees.

Summary

While there is increasing use of public cloud for strategic applications, archiving solely on the cloud would effectively delegate archive responsibility to the cloud vendor, which is a responsibility they may not accept. While there may be a benefit to strategically placing a copy of an archive on the cloud, retaining the primary archive copy at a government site is advisable. Costs would be a significant consideration, particularly if distributing from deep cloud storage such as Glacier. AWS Glacier was not assessed in the final evaluation. Transfer rates are difficult to judge because there is the question of whether data is being generated or copied from AWS or from another source coming across a Wide-Area-Network (WAN) connection, and throttled by that.

Tables

3.1 Design criteria

The design criteria and target market of a drive are interrelated. LTO8 is targeted to the backup market, as demonstrated by LTO marketing. The TS1155 is 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 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:

(((250 - serpentine passes) /10) +

(absolute value of error rate exponent / 2) +

(construction 3=moderate usage, 5=high usage))

/ 1.67 (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
IBM LTO8	6656/208	Backup	BaFe	$1x10^{-19}$	Moderate usage	Contact	10.0
IBM TS1155	7680/240	Enterprise	BaFe	$1x10^{-20}$	High usage	Contact	9.6

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 300 MB/sec, with 350 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 the previous generation. The sources 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
IBM LTO8	**360 MB/sec**	**Vendor**	**325 MB/sec**	**90.3%**	**329 MB/sec**	**91.3%**	**9.1**
IBM TS1155	**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 for this study is 8 TB, with 10 TB or more desired. Both reviewed technologies exceed the 8 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 is noted in Table 6. Capacity yield varies by media vendor.

Table 6. Storage Capacities.

Tape Drive Technology	Advertised Native Capacity	Actual/estimated Native Capacity	% of Advertised Capacity	Ranking
IBM LTO8	**12.0 TB**	**11.40 TB estimated**	**95.0% estimated**	**8.0**
IBM TS1155	**15.0 TB**	**14.25 TB estimated**	**95.0% estimated**	**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, IBM TS11xx 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[5]	Media $/TB	Ranking Drive Cost	Ranking Media Cost/TB
IBM LTO8	**$19,200**	**$151**	**$13.25**	**10.0**	**10.0**
IBM TS1155	**$20,000**	**$234**	**$16.42**	**9.6**	**8.1**

[5] Cited media costs are at quantity 20. Lower costs can be attained at higher quantities.

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 for media, a significant drop in media prices often occurs within 6 months after drive 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. Measured, or estimated, native capacities are used. Costs do not include system interfaces or cables, but do include 1 year of maintenance on the drives. Whole cartridges are costed, where only a partial cartridge is required in the final cartridge to reach the stated scenario capacity.

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 IBM TS11 drives have historically written to 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
IBM LTO8	**$41,118**	**$63,036**	**$90,088**	**10.0**
IBM TS1155	**$43,510**	**$66,786**	**$96,614**	**9.4**

3.6 Vendor Analyses

When selecting an archive technology, it makes sense to look at the company and product histories. An analysis of each company and the stability of each technology is shown in Table 9. For this revision of the study, IBM produces both of the alternative drives, so the comparison comes down to the longevity of LTO and TS11 product lines. IBM is an established and stable company; therefore, this rating should not be viewed as a market analysis. The longevity rankings were determined by the following formula:

(company age + technology age) / 12.5 (to adjust the highest rank to 10)

Table 9. Vendor Analyses

Company	Technology	Years in Business	Technology Age in Years	Longevity Ranking
IBM	LTO	107 (1911)	18 (2000)	10.0
IBM	3592	107 (1911)	15 (2003)	9.8

3.7 Drive Compatibility

The level of intergeneration drive read 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 current drive being evaluated. The following formula determined the ranking:

(Previous Generations Read + Future Generations Planned) x 2 (to adjust the highest rank to 10)

Table 10. Drive Compatibility

Technology	Previous Generations Read	Future Generations Mapped	Ranking
IBM LTO8	**1**	**4**	**10.0**
IBM TS1155	**2**	**3**	**10.0**

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
IBM LTO8	**10.0**	**8.0**	**10.0**	**10.0**	**9.1**	**10.0**	**10.0**	**10.0**
IBM TS1155	**9.6**	**10.0**	**8.1**	**10.0**	**10.0**	**9.6**	**9.8**	**9.4**

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

Selection Criteria	Wt	IBM LTO8	IBM TS1155	IBM LTO8	IBM TS1155
Design criteria		10.0	9.6	0.0	0.0
Capacity	20	8.0	10.0	160.0	200.0
Media cost/TB		10.0	8.1	0.0	0.0
Compatibility	15	10.0	10.0	150.0	150.0
Transfer rate	15	9.1	10.0	136.5	150.0
Drive cost		10.0	9.6	0.0	0.0
Vendor analyses	15	10.0	9.8	150.0	147.0
Scenario cost	35	10.0	9.4	350.0	329.0
Total Weighted Score				**946.5**	**976.0**

4.2 Conclusions and Notes

- TS1155 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 TS1155 would be incompatible with the existing Oracle SL8500 robotic library.
- As TS1155 and LTO8 drives were not available to be tested for this study, performance and capacity figures were based on vendor benchmarks, where available, or on drive specifications combined with predecessor 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. This will be a challenge when the Oracle T10000D drives are retired, since the only remaining drive technology still supported will be LTO, unless Oracle allows IBM TS11 attachment in the future.
- As any drive saturates the market, media costs drop, particularly for LTO since there are two cartridge manufacturers and LTO comprises over 96% of the tape market. With LTO7, tape composition switched from MP (Metal Particle) to the BaFe (Barium Ferrite) to attain higher storage densities. This change also has improved media stability.
- With proper handling and multiple copies, either 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.

- With Oracle ending the development and sale of its T10000 enterprise drives, only the IBM TS11 enterprise drives and multi-manufacturer LTO drives remain in the tape marketplace. The BER of the LTO7 was improved, TMR (Tunnel Magnetoresistance) tape head technology was implemented for LTO8, and the move to BaFe media has improved stability and density. These changes and others are blurring the lines between enterprise and non-enterprise drives. Quantum has proposed that LTO is now an Enterprise drive (Quantum, 2018).

4.3 Recommendations

Although the IBM TS1155 scored highest in this study, there is no compelling reason to adopt a new standard archive device at this time. The TS1155 would not be compatible with the existing Oracle SL8500 robotic tape library. Automation would require the purchase of another tape library, and substantial engineering effort would be required to integrate the two storage systems.

1. The USGS should proceed with plans to deploy LTO8 as the offline storage media of choice, then test and move to LTO9, when available.
2. Data archived on LTO6 should be regenerated on LTO8, as is planned for FY19.
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 LTO8, and store offline/offsite copies on LTO8. This strategy partially mitigates the risks of one or the other technology failing or being retired prematurely. This practice has been implemented and should continue. When LTO is eventually the only drive supported by the SL8500 tape library, a copy of the archive could exist on two different LTO generations of media.

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 continue. 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 implemented and should continue. Vendor verification features could be used, instead of or, in addition to file checksums.
7. All data should be migrated to new media approximately 5 years after writing. 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 for recovery and storage, and drives will become difficult to maintain. This practice has been implemented and should continue.
8. As archive media is retired at least 10% of the media should be read before disposal, to assess storage conditions and media viability. This practice has been implemented and should continue. During the recent read verification of 10% of the LTO5 media being disposed of, seven tapes failed to read on multiple drives, including one tape with a broken leader block. The number of verified tapes was increased when failures were noted. Several other tapes read slowly or failed on one drive but succeeded on another drive. All of the failed tapes had been shipped to and stored at NARA,

and the NARA rate of failed tapes to tested tapes was 3%. The tapes passed tests before being shipped and were spread across multiple boxes. No cause has been determined.

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. This is not an issue at this time, as the two current manufacturers of LTO media, FujiFilm and Sony, produce quality media that often rebranded. Though unlikely, this could become a concern if a new media manufacturer emerges.
11. The USGS should plan to update periodically 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, but this 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. Continual consideration of new archival technologies, including public-cloud and SSD is advised.

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/

OSTA. (2003). *Understanding CD-R & CD-W.* Retrieved from http://www.osta.org/technology/pdf/cdr_cdrw.pdf

Quantum. (2018, January). *LTO: The New "Enterprise Tape Drive".* Retrieved from Quantum Corporation: https://iq.quantum.com/exLink.asp?49663734OA24V33I185023536&WP00233A&view=1

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

Seppala, T. J. (2017, Aug 08). *Sony and IBM created 330TB data tapes for a massive analog archive*. Retrieved from engadget: https://www.engadget.com/2017/08/02/sony-ibm-330tb-data-storage-tapes/

Sony. (2017, August 2). *Sony Develops Magnetic Tape Storage Technology with the Industry's Highest*1 Recording Areal Density of 201 Gb/in2*. Retrieved from Sony.net: https://www.sony.net/SonyInfo/News/Press/201708/17-070E/index.html

Appendix: Supplemental Information

Vendor Sites

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/ltfs_supported_tape_drives_library_edition.html (IBM)

https://www.quantum.com/en/products/tape-storage/ (Quantum)

http://www.redbooks.ibm.com/redbooks/pdfs/sg245946.pdf (IBM)

https://iq.quantum.com/exLink.asp?54822614OR99B14I201949090 (Quantum PDF download)

https://www.sony.net/SonyInfo/News/Press/201708/17-070E/index.html (Sony)

Consortium Sites

http://www.lto.org/

Other

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

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

https://spectralogic.com/features/ts1155-technology-tape-drives

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

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

https://www.tapeandmedia.com/lto-8-tape-media-tapes.asp

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

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

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

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

https://phys.org/news/2017-06-chemical-solution-digital-storage.html

https://www.theregister.co.uk/2017/06/14/spectralogic_foresees_ibm_becoming_the_sole_tape_drive_supplier/

https://spectralogic.com/wp-content/uploads/white-paper-digital-data-storage-outlook-2017-v3.pdf

https://www.engadget.com/2018/03/19/nimbus-data-releases-record-100tb-ssd/

https://www.prnewswire.com/news-releases/fujifilm-corporation-receives-a-favorable-final-determination-in-united-states-itc-case-against-sony-corporation-300611522.html