Technical Support Services Contract (TSSC)					
USGS EROS	G10PD00044	Deliverable ID	TSSC-RSCEOSACP-TS-0001		
FY12 Offline Archive Media Trade Study		Date	08/09/12	Version Nbr	1.0

ARCHIVE AND RECORDS MANAGEMENT

FISCAL YEAR 2012 ARCHIVE MEDIA TRADE STUDY

STINGER GHAFFARIAN TECHNOLOGIES

ERT

Virtuoso
Evaluate. Integrate. Execute.

INFORMATION DYNAMICS

TSSC FISCAL YEAR 2012 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 their 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: Ken Boettcher, Ken Gacke, Cheryl Greenhagen, and Brian Reutzel.

[1] Stinger Ghaffarian Technologies, Inc., contractor to the U.S. Geological Survey, work performed under U.S. Geological Survey contract G10PC00044.

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

Contents

Figures

Tables

Abbreviations

AIT	Advanced Intelligent Tape
BER	Bit error rate
CD-ROM	Compact Disc - read only memory
CPU	Central processing unit
CRC	Cyclic redundancy check
DCT	Digital cassette tape

DLT	Digital Linear Tape
DVD	Digital Video Disc
EO	Erasable Optical
EROS	Earth Resources Observation and Science
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
MB	Megabytes (1,048,576 bytes)
NARA	National Archives and Records Administration
Q1, Q2, Q3, Q4	Fiscal or calendar quarter
QIC	Quarter-Inch Cartridge
SAIT	Super Advanced Intelligent Tape
SAM	Storage Archive Manager
SD	South Dakota

SDLT	Super Digital Linear 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 as long as 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 number 2 and 3.

June 2010

- Overall refresh of study, removing most references to older technologies (T10000, LTO4, DLT).
- Changed T10000B, LTO5, and TS1130 to be 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 be 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.

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 master and offsite copies where the working copy is stored nearline.

Linear Tape-Open 5 (LTO5) has been the archive media of choice at Earth Resources Observation and Science (EROS) for the past 18 months. LTO6 testing will begin in fiscal year 2013 (FY13). 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 have shown 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 T10000B/T10000C tape drives, Oracle LTO3/LTO4/LTO5 tape drives, an Oracle host server, Oracle Storage Archive Manager (SAM) HSM software, and a multi-vendor disk cache. The architecture of this HSM was determined by a trade study using a different set of requirements than this study.

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 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, in Sioux Falls, S. Dak., 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).

[HP, Hewlett-Packard; GB, gigabyte; MB, megabyte; MB/sec, megabyte per second]

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
3490	1995–2003	900 MB	2.7 MB/sec	Digital
DLT 7000	1996–2006	35 GB	5.0 MB/sec	Digital
DCT (Ampex DCRsI)	1992–2007	45 GB	12.0 MB/sec	Analog
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**	**1.5TB**	**140.0 MB/sec**	**Digital**
Oracle T10000C	**2012-present**	**5TB**	**240.0 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 do not need to be on similar media.
- 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 experience gained from benchmarking.

1.4 Selection criteria

The following criteria were used to determine which 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 at least 1 terabyte (TB) [1,000 gigabyte (GB)] capacity of uncompressed data.

3. The technology must have an uncompressed write transfer rate of at least 120 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.

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 unreliable in the past.

The following currently available drive technologies were selected for consideration:

1. Oracle T10000C
2. Oracle branded Hewlett-Packard (HP) LTO5 (Linear Tape Open)—representative of models by IBM, Quantum, and Tandberg.
3. IBM TS1140

The following future drive technologies are mentioned but not considered:

1. Oracle T10000D
2. HP LTO6
3. IBM TS1150

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 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 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, and performance. Serving frequently used working copies on disk provides significant performance benefits, although an offline master copy must be retained.

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.

1.5.3 Quantum DLT

In past revisions of this study, Quantum presented a viable challenge to LTO in the form of the Digital Linear Tape (DLT) line. DLT has lost substantial market share to the point that further development of the line has been discontinued (Betts, 2007). Although drives are still available, lack of further development has ensured that DLT is no longer competitive with LTO, and the specifications do not meet the minimum for this study. Quantum now produces LTO drives.

1.5.4 Tandberg/Exabyte VXA320, Sony SAIT-1/SAIT-2

Tandberg/Exabyte has evolved their early helical scan technology into the VXA320 with a native capacity of 160 GB and a native transfer rate of 12 MB/sec (Tandberg Data, 2012). This technology is based on consumer-grade cartridge and drive technologies. Although media costs are low, transfer rates are also low, and the USGS experience with consumer-grade storage technologies has shown that these technologies cannot withstand the rigors of a long-term archive.

Tape drives such as the 8 mm/Exabyte, which became popular in the 1990s, were based on consumer-grade helical scan technology and were notably slow and unreliable. Long start/stop/repositioning times dictated that if data were not kept streaming, the effective transfer rate dropped drastically. The necessarily complex drive path led to problems: 8 mm drives mangled tapes, and a confusing array of firmware versions often yielded unpredictable behavior and hangs. The transition from a market once ruled by 4 mm/8 mm helical scan drives to one ruled by LTO/DLT happened quickly, and the small current market share of helical scan technologies may indicate that the marketplace still remembers the difficulties of earlier helical scan drives. The market may never reconsider whether the earlier problems are overcome unless new terminology replaces "helical scan."

The Sony Super Advanced Intelligent Tape 1 (SAIT-1) and SAIT-2 seemed promising when first announced but were late to market, had slow transfer rates, and never gained sufficient market saturation to lower media costs. The SAIT-2 is reportedly only available in a Sony robotic library, which is targeted to video automation in the television industry.

1.5.5 DVD, Blu-Ray

The Digital Video Disc (DVD) and related technologies seem promising from the standpoint of expected longevity of the media; however, studies have shown that optical media can degrade and become unusable in as little as 5 years. Low capacity per media, low transfer rates, lack of media protection (no shell), no single standard, and high media costs add up to 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.

1.5.6 Newer optical 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.

Technical Assessment

2.1 Analysis

This technical assessment includes drives selected for final evaluation (T10000C, LTO5, TS1140) and drives anticipated to be released in the next two years (T10000D, LTO6, TS1150). LTO drives are available from multiple vendors (Tandberg, Quantum, IBM, HP), with an Oracle branded HP 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 T10000C**
- Oracle T10000D
- **HP LTO5**
- HP LTO6
- **IBM TS1140**
- IBM TS1150

Table 3. Technology comparison (yellow highlighted text indicates unverified information)

[TB, terabyte; MB/sec, megabyte per second; MB, megabyte; GB, gigabyte; est, estimated; TBD, to be determined; HW, hardware]

Specification	T10000C	T10000D	HP LTO5	HP LTO6	TS1140	TS1150
Uncompressed capacity	5.0 TB	8–10 TB	1.5 TB	3.2 TB	4.0 TB	8–10 TB
Uncompressed xfer rate	240 MB/sec	270–400MB/s	140 MB/sec	210 MB/sec	250 MB/sec	360 MB/sec
Recording technology	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine
Tracks	3,584		1,280		2,560	
Channels	32	32	16	16 or 32	32	32
Passes[2]	112		80		80	
Tape velocity (read)	5.62 m/s		6.1 m/s		3.17 m/s est[3]	
Type	Enterprise	Enterprise	Backup	Backup	Enterprise	Enterprise
Encryption support	HW option	HW option	HW built-in	HW built-in	HW built-in	HW built-in
Buffer size	2 GB	2 GB	256 MB	1GB	1 GB	1GB
Adaptive speeds	2	2	47–140 MB/s	Dynamic	13[4]	Multiple
Price (typical street)	$24,000 est	$24,000 est	$17,000	$17,000 est	$29,000 est	$29,000 est
Prev generations read	2	TBD	2	2	3	TBD
Prev generations written	0	TBD	1	1	1	TBD
Bit Error Rate (BER)	$1x10^{-19}$	$1x10^{-19}$	$1x10^{-17}$	$1x10^{-17}$	$1x10^{-20}$	$1x10^{-20}$
Drive manufacturers	1	1	4+	4+	1	1
Availability	2011	2013	2010	2012[5]	2011	2013?

[2] As reported by vendor or calculated by dividing tracks by channels.
[3] Estimated using transfer rate, bits per inch, and number of channels.
[4] Variously reported by IBM as 13 and as 14.
[5] Per Spectra-logic, via Storage Newsletter.

Oracle T10000C

The T10000C is the Oracle flagship high-capacity enterprise tape drive typically used in conjunction with Oracle robotic libraries, such as the SL8500. EROS has five T10000C drives for use in the SL8500, in addition to ten second-generation T10000B drives.

Advantages

- The T10000C is an evolution of the T10000/T10000B, which have performed reliably for the USGS.
- Native capacity is 5 TB and native transfer rate is 240 MB/sec. The T10000C also can stream at a lower rate, which is important because most disks will not be able to keep up at 240 MB/sec. Reconfiguration or replacement of disks will improve their transfer rates.
- The T10000C uses 32 channels per pass (as compared to 16 on LTO), which increases the transfer rate.
- The T10000C is targeted to the enterprise storage market where data viability, speed, and capacity are more important than cost.
- The T10000C 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.
- T10000C drives provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- T10000C uses a new media design that is incompatible with previous drives, but it will allow media reuse with the T10000D.
- The T10000C has a 2 GB buffer, which prevents occasional data starvation from reducing the transfer rate.

- The Bit Error Rate (BER) is $1x10^{-19}$.
- A hardware encryption option module is available.
- Internal CRC ensures that there was 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.

Disadvantages

- Repeated 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). This repeated use should not be a concern for archive operations because usage is limited.
- The T10000C drives cost 41 percent more than LTO5 drives but cost 17 percent less than the TS1140.
- Based on sales of the T10000/T10000B, the T10000C sales are anticipated to be primarily for use in Oracle robotics. For this reason, the T10000C is anticipated to have a market share that will remain low compared to LTO, ensuring that media costs will remain high.
- The T10000C drive is only available from Oracle. This limited availability keeps the price high but does eliminate concerns of incompatibility.

Summary

The T10000C 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 T10000B or T10000C, use of an enterprise solution for the second copy is not warranted.

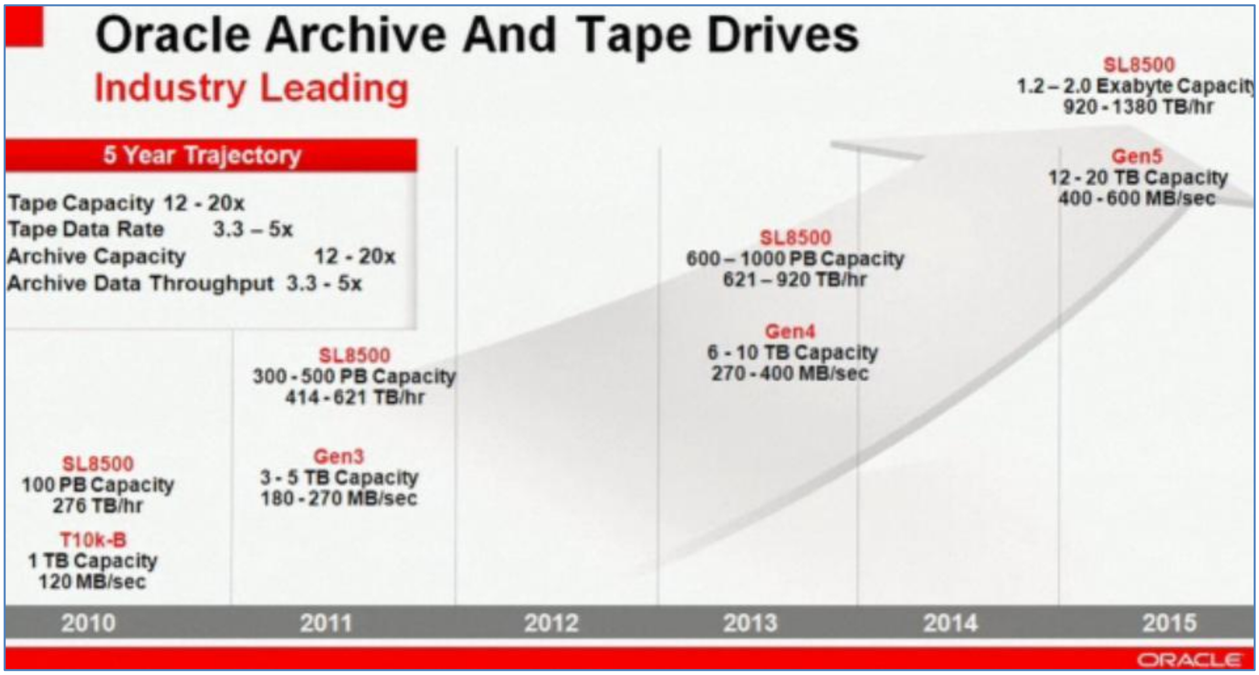

Figure 1. Oracle Roadmap (uncompressed) (Oracle, 2010).

Oracle T10000D

The T10000D is the fourth generation of the T10000 line. The T10000D is anticipated to ship in 2013 per Oracle.

Advantages

- The T10000D is an evolution of the T10000/T10000B/T10000C, which have performed reliably for the USGS.
- Native capacity of 8–10 TB and native transfer rate of 270–400 MB/sec. The T10000D is expected to stream at lower rates, which is important because most disks will not be able to keep up at 270+ MB/sec.
- The T10000D is expected to use at least 32 channels per pass (compared to 16 on competing drives), which increases the transfer rate. It is possible that Oracle will raise the number of channels to 64 to increase the transfer rate and reduce serpentine passes even if the number of tracks is increased to add capacity.
- The T10000D is targeted to the enterprise storage market where data viability, speed, and capacity are more important than cost.
- The media for the T10000D is expected to be the same as the T10000C, and the T10000D is expected to be able to read media written on the T10000B/T10000C, and possibly the T10000. The T10000C/T10000D media 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 expected to be compatible with the SL8500.
- As with predecessor drives, T10000D drives should provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.

- Some future follow-on drives may use the same media, which would allow savings through media reuse.
- The T10000D 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 $1x10^{-19}$ or better.
- A hardware encryption option module is anticipated.

Disadvantages

- The T10000D drives are expected to cost more than the LTO but less than the TS1140.
- Based on sales of predecessor drives, the T10000D 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 T10000D drive is expected to be available only from Oracle. This availability keeps the price high but does eliminate concerns of incompatibility.

Summary

The T10000D should replace the T10000C drive as the flagship high-capacity enterprise drive typically used in conjunction with Oracle robotic libraries because the T10000D should be priced comparably. The T10000D is not yet available and was therefore not assessed in the final evaluation.

HP LTO5

The LTO5 is the most recent available generation of the LTO tape family. EROS has four LTO5 drives for use in the SL8500, in addition to several LTO4 and LTO3 drives. Several non-robotic LTO drives are also in use at EROS.

Advantages

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of Q1 2012, LTO held an 89.9 percent market share (Santa Clara Consulting Group, 2012). Several competing technologies, such as DLT (Betts, 2007) and SAIT, have been driven from the market.
- Native capacity is 1.5 TB and native transfer rate is 140 MB/sec.
- The HP LTO5 drive can adapt the transfer rate to match the streaming speed of the source.
- LTO5 is backward read compatible with LTO3 and LTO4, and backward write compatible with LTO4 (at the lower LTO4 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.
- The LTO5 has a 256 MB buffer that prevents occasional data starvation from reducing the transfer rate.
- Hardware encryption is available, as shown in figure 2.

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.

- Repeated end-to-end use of a tape would be a concern because one end-to-end read/write incurs 80 passes (1,280 tracks divided by 16 channels). This repeated use 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 substantially 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.
- LTO was co-developed by IBM, HP, and Quantum (acquired from Seagate/Certance). 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 has observed that drives slowly degrade without failing, 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

Testing of LTO5 technology at EROS was completed in early 2011. EROS has encountered occasional issues with LTO5 where there is no error on write, and it is not detected until data are read on non-robotic drives. This seems to be limited to a specific brand of media.

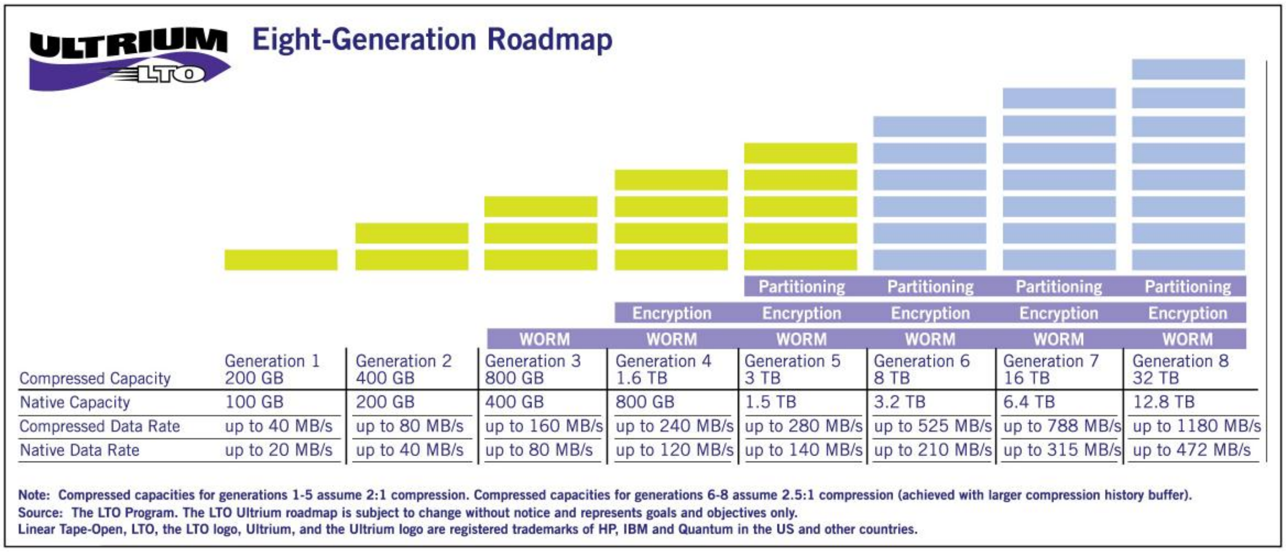

Figure 2. LTO Roadmap (compressed/uncompressed) (source: LTO Consortium).

HP LTO6

The LTO6 is the next anticipated generation of the LTO tape family, with release expected in late 2012 based on a typical LTO release cycle of 2 years, and on comments made by vendors.

Advantages

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of Q1 2012, LTO held an 89.9 percent market share (Santa Clara Consulting Group, 2012). Several competing technologies, such as DLT (Betts, 2007) and SAIT, have been driven from the market.
- Native capacity is expected to be 3.2 TB and native transfer rate is expected to be 210 MB/sec.
- The HP LTO6 drive is anticipated to use an adaptive transfer rate to match the streaming speed of the source.
- LTO6 should be backward read compatible with LTO4 and LTO5, and backward write compatible with LTO5 (at the lower LTO5 capacity).
- 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.
- Hardware encryption is anticipated.

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.

- Repeated end-to-end use of a tape would normally be a concern because one end-to-end read/write is expected to incur 80 or more passes. This repeated use should not be a concern for archive operations because offline tapes are typically written once and read once during verification.
- Each generation of LTO requires new media to attain the rated capacity, ensuring that media costs will be substantially 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 was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant use.
- LTO was co-developed by IBM, HP, and Quantum (acquired from Seagate/Certance). 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

LTO6 is expected to be announced in mid-2012 and made available in late 2012. LTO6 is not yet available and was not assessed in the final evaluation.

IBM TS1140

The TS1140 is an enterprise-class tape drive, used primarily in IBM robotic libraries and autoloaders. The TS1140 is a follow-on drive to the TS1130.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, and TS1130.
- Supports dual 8 gigabit per second (Gbit/sec) Fiber Channel interfaces.
- Native capacity is 4 TB and native transfer rate is 250 MB/sec.
- The TS1140 is a robust storage technology, with the tape cartridge and drive built to withstand constant or frequent use in a robotic environment.
- The TS1140 uses some of the same media as earlier generations, plus a new higher capacity cartridge.
- A hardware encryption feature is included in the drive, as shown in figure 3.

Disadvantages

- Designed primarily for use in IBM robotic libraries.
- Not supported by the Oracle SL-8500 robotic tape library.

Summary

Enterprise-class robustness is not required when the working copy of a dataset is already on enterprise-class technology in the EROS robotic library. IBM recently reported development of a recording method that will yield a capacity of 35 TB per cartridge (Mearian, 2010) but did not reveal a timeline.

3592 Model	Gen 1 3592 J1A	Gen 2 TS1120	Gen 3 TS1130	Gen 4 TS1140	Gen 5	Gen 6
Shipped	3Q2003	3Q2006	3Q2008	2Q2011		
Max Cartridge Capacity	300 GB	700 GB	1.0 TB	4 TB	8-10 TB	14-20 TB
Native capacity	300 GB JA	500 GB JA 700 GB JB	640 GB JA 1.0 TB JB	1.6 TB JB 4 TB JC	6-8 TB JC 8-10 TB JD	6-8 TB JC 14-20 TB JD
Data Rate MB/S	40	100	160	250	Up to 360	Up to 540
Cartridge Support JA type JB type JC type JD type	JA/JJ/JW/JR	JA/JJ/JW/JR JB/JX	JA/JJ/JW/JR JB/JX	JA/JJ/JW/JR JB/JX JC/JY/JK	JC/JY/JK JD/JZ	JC/JY/JK JD/JZ
Encryption	N/A	Yes	Yes	Yes	Yes	Yes
Partitioning / LTFS Support	N/A	N/A	N/A	Yes	Yes	Yes
Server Attachment	Fibre FICON ESCON	Fibre FICON ESCON	Fibre FICON ESCON	Fibre FICON ESCON	Fibre FICON FCoE	Fibre FICON FCoE

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

IBM TS1150

The TS1150 is anticipated to be the next generation of the 3592 tape family, with release expected in 2013. Note that the TS1150 name has not been confirmed, but follows logically.

Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, TS1130, and TS1140.
- Should support at least two 8 Gbit/sec Fiber Channel interfaces.
- Native capacity is expected to be 8–10 TB and native transfer rate may reach 360 MB/sec.
- The TS1150 will be a robust storage technology, with the tape cartridge and drive built to withstand constant or frequent use in a robotic environment.
- The TS1150 may use some of the same media as the TS1140.
- A hardware encryption feature should be included in the drive.

Disadvantages

- Designed primarily for use in IBM robotic libraries.

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. TS1150 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. LTO5 is targeted to the backup market, as demonstrated by LTO marketing. The T10000C and TS1140 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 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 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 formula used to rank design criteria was:

((150-serpentine passes)/10)+

(absolute value of error rate exponent/2)+

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

(head contact 3=contact, 5=min contact)

/ 2.5 (to adjust the highest rank to 10)

Table 4. Design criteria and target market. [TBD, to be determined; MP, metal particle]

Technology	Serpentine tracks/ Passes	Target market	Tape composition	Uncorrected error rate	Cartridge construction rating	Head contact	Ranking
Oracle T10000C	3584/112	Enterprise	Advanced MP	$1x10^{-19}$	High usage	Min contact	9.3
HP LTO5	1280/80	Backup	Thin film MP	$1x10^{-17}$	Moderate usage	Contact	8.6
IBM TS1140	2560/80	Enterprise	Barium Ferrite	$1x10^{-20}$	High usage	Contact	10.0

3.2 Transfer Rate

Transfer rate is important because it establishes how quickly the migration of an archive dataset may be completed and how fast a production system may generate products from the archive media. The minimum transfer rate requirement is 120 MB/sec, with 140 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 test results also applies to capacities 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. [%, percent, MB/sec, megabyte per second]

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 adv.	Rank ing
Oracle T10000C	240 MB/sec	EROS testing	200 MB/sec	83.3%	236 MB/sec	98.3%	8.7
HP LTO5	140 MB/sec	EROS testing	119 MB/sec	85.0%	124 MB/sec	88.6%	4.8
IBM TS1140	252 MB/sec	Vendor	252.0 MB/sec	100%	250.0 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 LTO5 at 1.41 TB of usable capacity per tape. The minimum capacity requirement is 1 TB, with 1.5 TB or more desired. All the reviewed technologies meet the 1 TB requirement based on the advertised capacity. Because much of the data archived is 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 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. [TB, terabyte; GB, gigabyte, % percent]

Tape drive technology	Advertised/ proposed native capacity	Measured/approximate native capacity	% of advertised capacity	Ranking
Oracle T10000C	5.0 TB	4.68 TB EROS measured	93.6%	10.0
HP LTO5	1.5 TB	1.41 TB EROS measured[6]	94.0%	3.0
IBM TS1140	4.0 TB	3.80 TB approximate	95.0% approximate	8.1

[6] EROS has experienced a highly variable capacity with LTO5, as low as 1.1TB. This may be because of inconsistent media manufacturing.

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 cheapest (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.

Maintenance could be a consideration but was removed from a previous iteration of the study because of the tenuous status of support costs and incomplete information. Comparing maintenance contracts can be complex because of terms that vary by vendor and because of the existence of third-party support.

It should be noted that unlike LTO, Oracle and IBM drives have 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. This would make a case for these technologies to be used for nearline or onsite offline copies, because the media would be readily available for rewriting at higher capacity.

Table 7. Drive and media costs. [$/each, dollars per each; $/TB, dollars per terabyte; est, estimated]

Tape drive technology	Drive $/each	Media $/each	Media $/TB	Ranking drive cost	Ranking media cost/TB
Oracle T10000C	$24,000	$249	$50	7.1	5.8
HP LTO5	$17,000	$44	$29	10.0	10.0
IBM TS1140	$29,000	$232	$58	5.9	5.0

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 competition, a considerable drop in media prices often occurs within 6 months after product introduction.

Rankings are based on the 100 TB option and were established by setting the cheapest to 10, and then rating each of the others against the lowest cost. Advertised native capacities are used. Costs do not include maintenance, system interfaces, or cables.

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 typically have used the same media for at least two generations.

Table 8. Scenario costs (drives, media) [TB, terabyte]

Technology	100 TB 2 drives	200 TB 3 drives	400 TB 4 drives	100 TB ranking
Oracle T10000C	$53,000	$82,000	$116,000	7.0
HP LTO5	$36,900	$56,800	$79,600	10.0
IBM TS1140	$58,000	$87,000	$116,000	6.4

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 T10000C, the technology was based on the ancestor 9940; therefore, the technology age includes the 9940. The longevity rankings were determined by the following formula:

(company age + technology age) / 11.8 (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	43 (1969)	12 (2000)	4.7
HP	LTO	73 (1939)	12 (2000)	7.2
IBM	3592 (3590)	101 (1911)	17 (1995)	10.0

3.7 Drive compatibility

The level of intergeneration drive compatibility and planned future drives are shown in table 10. The columns "% Previous generations read" and "% Previous generations written" indicate the percentage of previous generations that are read/written by the generation indicated. Drives that are the first generation receive a score of 50 percent, so the first generation product will not be penalized. The column "Future generations planned" indicates the number of generations planned in the current drive family, following the drive indicated. The ranking was determined by the following formula:

(% Previous generations read + % Previous generations written + (Future generations planned x 20)) / 19 (to adjust the highest rank to 10)

Table 10. Drive compatibility (yellow highlighted text indicates unverified information) [%, percent]

Technology	% Previous generations read	% Previous generations written	Future generations planned	Ranking
Oracle T10000C	100	0	2	7.4
HP LTO5	50	25	3	7.1
IBM TS1140	100	50	2	10.0

3.8 Ranking summary

The ranking summary provides a quick reference to the rankings.

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

Drive	Design criteria	Capacity	Media cost	Drive compatibility	Transfer rate	Drive cost	Vendor analyses	Scenario cost
T10000C	9.3	10.0	5.8	7.4	8.7	7.1	4.7	7.0
HP LTO5	8.6	3.0	10.0	7.1	4.8	10.0	7.2	10.0
IBM TS1140	10.0	8.1	5.0	10.0	10.0	5.9	10.0	6.4

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 and write 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. (See table 8)

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

Table 12. Weighted decision matrix [wt, weight; /TB, per terabyte]

Selecton criteria	Wt	Oracle T10000C	HP LTO5	IBM TS1140	Oracle T10000C	HP LTO5	IBM TS1140
Design criteria		9.3	8.6	10.0	0.0	0.0	0.0
Capacity	20	10.0	3.0	8.1	200.0	60.0	162.0
Media cost/TB		5.8	10.0	5.0	0.0	0.0	0.0
Compatibility	15	7.4	7.1	10.0	111.0	106.5	150.0
Transfer rate	15	8.7	4.8	10.0	130.5	72.0	150.0
Drive cost		7.1	10.0	5.9	0.0	0.0	0.0
Vendor analyses	15	4.7	7.2	10.0	70.5	108.0	150.0
Scenario cost	35	7.0	10.0	6.4	245.0	350.0	224.0
Total Weighted Score					**757.0**	**696.5**	**836.0**

4.2 Conclusions and notes

TS1140 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 scores.

A TS1140 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 T10000C, as are most archives at EROS.
- 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 T10000C and TS1140 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 TS1140 scored highest in this study, there is no compelling reason to adopt a new standard archive device at this time. The TS1140 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 LTO5 as the offline storage media of choice, then test and move to LTO6 when available. T10000C recently has been used for some onsite offline copies, leveraging Oracle plans to utilize the same media for T10000C and T10000D.

2. Data archived on LTO2, LTO3, and LTO4 should be migrated to LTO5 or LTO6 in the next 2 years.
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 T10000C, store offline/onsite data on LTO5 or T10000C, and store offline/offsite copies on LTO5. 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 the National Archives and Records Administration (NARA), but should be frequent enough to provide an opportunity to detect deteriorating media.
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 partially implemented and should continue.
7. All data should be migrated to new media 3 to 5 years after it was 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. 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.

9. Where possible, the USGS should avoid buying media brands that have proven unreliable.
10. 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. There currently is no compelling reason to abandon LTO technology.
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.

References Cited

Betts, B., 2007, DLT tape reaching end of reel: The Register, available at *http://www.theregister.co.uk/2007/03/21/lto_beats_dlt/*. (Accessed August 9, 2012.)

Mearian, L., 2010, Fujifilm, IBM tech could yield 35TB magnetic tape drives: Computerworld, available at *http://www.computerworld.com/s/article/9147498/Fujifilm_IBM_tech_could_yield_35TB_magnetic_tape_drives*. (Accessed August 9, 2012.)

Oracle, 2010, Oracle-Sun 08/2010 Systems Strategy: Network Management, available at *http://netmgt.blogspot.com/2010/08/oracle-sun-082010-webcast-highlights.html*. (Accessed August 9, 2012.)

Santa Clara Consulting Group, 2012, LTO Accounted for 90% of WW Tape Market in 1Q12: Micro-Journal, available at *http://www.storagenewsletter.com/news/marketreport/santa-clara-consulting-group-tape-1q12*. (Accessed August 9, 2012.)

Tandberg Data, 2012, VXA-320—Product overview: Tandberg Data, available at *http://www.tandbergdata.com/us/?linkServID=701A7BFD-A8ED-0BB1-B935800DD905FADF*. (Accessed August 9, 2012.)

Appendix: Supplemental Information

Vendor sites

http://h18006.www1.hp.com/storage/tapestorage/tapedrives.html (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)

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

http://www.tandbergdata.com/us/en/products/drives/lto/ (Tandberg)

Consortium sites

http://www.lto.org/newsite/index.html

Other

http://ieeexplore.ieee.org/application/enterprise/entconfirmation.jsp?arnumber=1065475

http://netmgt.blogspot.com/2010/08/oracle-sun-082010-webcast-highlights.html

http://www.tapeandmedia.com/lto-6-tape-media-tapes.asp

http://www.storagenewsletter.com/news/tapes/first-step-for-lto-6

http://www.storagenewsletter.com/news/tapes/spectra-logic-lto-6-to-ship-2h12

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

http://technorati.com/technology/article/sony-signs-major-licensing-deal-to/

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.theregister.co.uk/2010/08/11/oracle_tape_roadmap/

http://www-03.ibm.com/systems/storage/tape/ts1140/index.html

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/82f67152325e844985257960005866fa/$FILE/IBM%20TS1140%20Technology%20White%20Paper%2011%20October%202011%20Final%20v3.pdf

http://www.clipper.com/research/TCG2004040.pdf

http://www.redbooks.ibm.com/redbooks/pdfs/sg244632.pdf

http://www.computerworld.com/hardwaretopics/storage/story/0,10801,110667,00.html?source=NLT_SU&nid=110667

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

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

http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA0-7675ENW.pdf