

# **Archive and Records Management—Fiscal Year 2010 Offline Archive Media Trade Study**

Open-File Report 2010–1222



# **Archive and Records Management—Fiscal Year 2010 Offline Archive Media Trade Study**

By Tom Bodoh, Ken Boettcher, Ken Gacke, Cheryl Greenhagen, and  
Al Engelbrecht

Open-File Report 2010–1222

**U.S. Department of the Interior**  
**U.S. Geological Survey**

**U.S. Department of the Interior**  
KEN SALAZAR, Secretary

**U.S. Geological Survey**  
Marcia K. McNutt, Director

U.S. Geological Survey, Reston, Virginia: 2010

For more information on the USGS—the Federal source for science about the Earth, its natural and living resources, natural hazards, and the environment, visit <http://www.usgs.gov> or call 1-888-ASK-USGS

For an overview of USGS information products, including maps, imagery, and publications, visit <http://www.usgs.gov/pubprod>

To order this and other USGS information products, visit <http://store.usgs.gov>

Any use of trade, product, or firm names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

Although this report is in the public domain, permission must be secured from the individual copyright owners to reproduce any copyrighted materials contained within this report.

Suggested citation:

Bodoh, Tom, Boettcher, Ken, Gacke, Ken, Greenhagen, Cheryl, and Engelbrecht, Al, 2010, Archive and records management—Fiscal year 2010 offline archive media trade study: U.S. Geological Survey Open-File Report 2010–1222, 15 p. plus app.

## **Preface**

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



# Contents

Preface .....	iii
Abstract .....	1
Revision History.....	1
February 2004.....	1
September 2006.....	1
June 2008.....	1
June 2010.....	2
Introduction.....	2
Purpose and Scope .....	2
Background.....	2
Data Integrity.....	2
Selection Criteria .....	3
Dismissed Technologies .....	4
Magnetic Disk.....	4
Solid State Disk (SSD).....	4
CD-ROM, DLT 8000, QIC, Mammoth, and Erasable Optical (EO) .....	4
Oracle 9840.....	4
Quantum DLT.....	4
Tandberg/Exabyte VXA320, Sony SAIT-1/SAIT-2 .....	4
DVD, HD-DVD, Blu-Ray .....	5
Newer Optical Technologies.....	5
Technical Assessment.....	5
Analysis .....	5
Oracle T10000B.....	6
Advantages.....	6
Disadvantages.....	6
Summary.....	6
Oracle T10000C.....	7
Advantages .....	7
Disadvantages.....	7
Summary.....	7
HP LT05.....	7
Advantages.....	7
Disadvantages.....	7
Summary.....	8
HP LT06.....	8
Advantages.....	8
Disadvantages.....	9
Summary.....	9
IBM TS1130 .....	9
Advantages .....	9
Disadvantages.....	9
Summary.....	9

IBM TS1140 .....	10
Advantages .....	10
Disadvantages .....	10
Summary .....	10
Tables .....	10
Design Criteria .....	10
Transfer Rate .....	10
Capacity .....	11
Cost Analysis .....	11
3.5 Scenarios .....	12
Vendor Analyses .....	12
Drive Compatibility .....	12
Ranking Summary .....	13
Conclusions and Recommendations for USGS Offline Archiving Requirements .....	13
Weighted Decision Matrix .....	13
Conclusions and Notes .....	14
Recommendations .....	14
Appendix: Citations .....	17
Vendor Sites .....	17
Consortium Sites .....	17
Other .....	17

## Figures

1. Screen capture showing the Oracle roadmap (uncompressed) .....	6
2. Screen capture showing the LTO roadmap (with 2:1 compression) .....	8
3. Screen capture showing the IBM roadmap (uncompressed) .....	9

## Tables

1. Recent and current archive technologies used at the U.S. Geological Survey .....	2
2. Tape drive markets and characteristics .....	3
3. Technology comparison .....	5
4. Design criteria and target market .....	10
5. Transfer rates .....	11
6. Storage capacities .....	11
8. Scenario costs (drives, media) .....	12
7. Drive and media costs .....	12
9. Vendor analyses .....	12
10. Drive compatibility .....	13
11. Ranking summaries .....	13

## Abbreviations and Acronyms

AIT	Advanced Intelligent Tape
BER	Bit Error Rate
CD	Compact Disc
CD-ROM	Compact Disc - Read Only Memory
CERN	Conseil European pour la Recherche Nucleaire
CRC	Cyclic Redundancy Check
CPU	Central Processing Unit
DCT	Digital Cassette tape
DLT	Digital Linear Tape
DVD	Digital Video Disc
EO	Erasable Optical
EOT	End of tape
EROS	Earth Resources Observation and Science
FYyy	Fiscal Year yy
GB	Gigabytes (1,024 MB, or 1,073,741,824 bytes)
Gbit/sec	Gigabit per second
HD-DVD	High Definition Digital Versatile Disc (formerly Digital Video Disc)
HDT	High Density Tape
HP	Hewlett-Packard
HVD	Holographic Versatile Disc
HW	Hardware
IBM	International Business Machines
LP DAAC	Land Processes Distributed Active Archive Center
LTO	Linear Tape-Open
MB	Megabytes (1,048,576 bytes)
MB/sec	Megabytes per second
NARA	National Archives and Records Administration
QIC	Quarter-inch Cartridge
SAIT	Super Advanced Intelligent Tape
SDLT	Super Digital Linear Tape
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	U.S. Geological Survey



# Archive and Records Management—Fiscal Year 2010 Offline Archive Media Trade Study

By Tom Bodoh<sup>1</sup>, Ken Boettcher<sup>1</sup>, Ken Gacke<sup>1</sup>, Cheryl Greenhagen<sup>1</sup>, and Al Engelbrecht<sup>1</sup>

## 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 (USGS). 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 2010 (FY10) revision of a study completed in FY01 and revised in FY03, FY04, FY06, and FY08.

## 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 Linear Tape-Open (LTO) 2 as a current archive technology.
- Added Super Advanced Intelligent Tape (SAIT)-1 and Super Digital Linear Tape (SDLT) 600 as considered drives.
- Replaced International Business Machines (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 showing 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 SAIT-2 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, SAIT-1, and SAIT-2.
- Modified so that future technologies are no longer scored.
- Decreased the number of drives for scenarios #2 and #3.

<sup>1</sup>Stinger Ghaffarian Technologies, Inc., contractor to the U.S. Geological Survey, work performed under USGS contract 08HQC0005.

## 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 due to lack of data.
- Adjusted minimum transfer rate and capacity to be considered for the study.

## Introduction

### 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 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.

Note that LTO4 has been the archive media of choice at USGS for the past 2 years. LTO5 testing will begin in FY10. There is no compelling reason for the USGS to change technologies away from LTO 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 eventually no longer meets USGS 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 USGS. The primary nearline mass-storage system at the Earth Resources Observation and Science (EROS) Center contains an HSM using an Oracle SL8500 robotic tape library, Oracle T10000/T10000B tape drives, Oracle LTO3/LTO4 tape

drives, an Oracle host server, Oracle SAM HSM software, and a multivendor 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 be used at the EROS Center and meeting USGS criteria. 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.

## Background

The USGS EROS Center, in Sioux Falls, South Dakota, has archived offline datasets using several technologies (table 1).

In 2003, the USGS migrated more than 50,000 3480 and 3490 tapes to nearline storage and to 110 LTO2 tapes. This migration was performed over a period of 5.5 months, slowed by the handling of the large number of 3480/3490 tapes. High Density Tape (HDT), 3480/3490, and Digital Cassette Tape (DCT) were proven to be robust and high-performance for their time. As technology advances, as datasets grow, as media ages, and as USGS Digital Library space fills, the USGS must migrate data to newer, more physically compact, and higher performing storage technologies.

## Data Integrity

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

**Table 1.** Recent and current archive technologies used at the U.S. Geological Survey.

[Current in bold. USGS, U.S. Geological Survey; HDT, High Density Tape; GB, gigabyte; MB/sec, megabyte per second; MB, megabyte; DLT, Digital Linear Tape; DCT, Digital Cassette Tape; HP, Hewlett-Packard; LTO, Linear Tape-Open]

Tape drive technology	Years used at USGS	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
<b>Oracle 9940B</b>	<b>2002–present</b>	<b>200 GB</b>	<b>30.0 MB/sec</b>	<b>Digital</b>
<b>HP LTO Ultrium 2</b>	<b>2003–present</b>	<b>200 GB</b>	<b>40.0 MB/sec</b>	<b>Digital</b>
<b>HP LTO Ultrium 3</b>	<b>2005–present</b>	<b>400 GB</b>	<b>80.0 MB/sec</b>	<b>Digital</b>
<b>HP LTO Ultrium 4</b>	<b>2007–present</b>	<b>800 GB</b>	<b>120.0 MB/sec</b>	<b>Digital</b>

- The number of archival copies: USGS archives must have working and master copies, and an offsite copy is desirable. The master and working copies need not 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 media are stored in a secure and climate-controlled environment.
- The composition of the media: Some media compositions last significantly 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 (table 2).
  - 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 major 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	As much as 100 percent 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 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.

Experience with 3480, 3490, 9840, 9940, and T10000 has shown Oracle/Sun/StorageTec (STK) products to be reliable, but the Oracle/Sun D3 helical scan drive was problematic and was discontinued quickly. On several occasions tapes that had unrecoverable errors were sent to Oracle for recovery. Some tapes were recovered, but some were unrecoverable because of cartridge contamination. Tape drive failures typically happen without tape damage and are replaced without causing data loss.

## Selection Criteria

The following criteria were used in determining which technologies should be considered.

1. The technology must be currently available and the most recent drive manufactured to be considered in the final analysis. 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 Archive and Records Management—Fiscal Year 2010 Offline Archive Media Trade Study

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.
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, DAT, and D3 have proven unreliable in the past.

The following currently available drive technologies were selected for consideration.

1. Oracle T10000B;
2. Hewlett-Packard (HP) LTO5 (Linear Tape Open)—representative of models by IBM, Quantum, and Tandberg; and
3. IBM TS1130.

The following future drives technologies are mentioned but not considered:

1. Oracle T10000C;
2. HP LTO6; and
3. IBM TS1140.

### Dismissed Technologies

The following technologies were dismissed from analysis or consideration.

### Magnetic Disk

Disk prices continue to drop, while 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 stay viable up to 10 years, the more costly disk is typically 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.

### Solid State Disk (SSD)

Similar to magnetic disk, SSD prices continue to drop, while reliability, performance, and capacity increase. It is expected that SSD, over 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.

### CD-ROM, DLT 8000, QIC, Mammoth, and Erasable Optical (EO)

This category includes technologies that are low capacity, low performance, or aged. All of these products have been available for some time but can immediately be dismissed on the basis of obvious limitations in performance, capacity, or reliability. These products are not a good fit for large digital archives.

### Oracle 9840

The Oracle 9840 is a fast-access technology used almost exclusively in conjunction with Oracle robotic libraries. Although it is an enterprise-class drive, it has low capacity, low transfer rate, and high cost. The advantage of this drive is the fast access: the dual reel design does not require a lengthy loading sequence, and it is positioned at tape midpoint for faster access. Although this technology is useful where fast nearline access is required, the technology offers minimal benefit in the offline archive media arena.

### Quantum DLT

In past revisions of this study, Quantum presented a viable challenge to LTO in the form of the DLT line. DLT has lost substantial market share to the point that further development of the line has been officially discontinued. 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.

### 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 24 MB/sec. 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 SAIT-1 and SAIT-2 seemed promising when first announced but were late to market, have 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.

### DVD, HD-DVD, Blu-Ray

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 would not work for high volume archival use.

High Definition Digital Versatile Disc (HD-DVD) was withdrawn from the marketplace after failing to compete with Blu-Ray. Blu-Ray would certainly have some application in distribution and short-term storage of large amounts of data, but like Compact Disc (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.

### 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. Rivals claim as much as 100 TB per disk will be possible.

## Technical Assessment

### Analysis

This technical assessment includes drives selected for final evaluation (T10000B, LTO5, TS1130) and drives anticipated to be released in the near future (T10000C, LTO6, TS1140) (table 3). LTO drives are available from multiple vendors (Tandberg, Quantum, IBM, HP), with HP 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 final evaluation.

- **Oracle T10000B**
- Oracle T10000C
- **HP LTO5**
- HP LTO6
- **IBM TS1130**
- IBM TS1140

**Table 3.** Technology comparison.

[Yellow highlighted text indicates unverified information. HP, Hewlett-Packard; LTO, Linear Tape-Open TB, terabyte; MB/sec, megabyte per second; TBD, to be determined; m/sec, meters per second; HW, hardware; MB, megabyte; GB, gigabyte; est, estimated]

Specification	T10000B	T10000C	HP LTO5	HP LTO6	TS1130	TS1140
Uncompressed capacity	1.0 TB	<b>2.0 TB</b>	1.5 TB	<b>3.2 TB</b>	1.0 TB	<b>2.0 TB</b>
Uncompressed xfer rate	120 MB/sec	<b>200+ MB/sec</b>	140 MB/sec	<b>210 MB/sec</b>	160 MB/sec	<b>240 MB/sec</b>
Recording technology	Serpentine	<b>Serpentine</b>	Serpentine	<b>Serpentine</b>	Serpentine	<b>Serpentine</b>
Tracks	1,152	<b>TBD</b>	1,280	<b>TBD</b>	1152	<b>TBD</b>
Channels	32	<b>32</b>	16	<b>TBD</b>	16	<b>TBD</b>
Passes	36	<b>TBD</b>	80	<b>TBD</b>	72	<b>TBD</b>
Tape velocity	3.74 m/sec	<b>TBD</b>	TBD	<b>TBD</b>	8.6 m/sec	<b>TBD</b>
Type	Enterprise	<b>Enterprise</b>	Backup	<b>Backup</b>	Enterprise	<b>Enterprise</b>
Encryption support	HW option	<b>HW option</b>	HW built-in	<b>HW built-in</b>	HW built-in	<b>HW built-in</b>
Buffer size	256 MB	<b>256 MB</b>	256 MB	<b>TBD</b>	1 GB	<b>1 GB</b>
Adaptive speeds	2 speeds	<b>2 speeds</b>	47–140 MB/s	<b>Dynamic</b>	6 speeds	<b>Multiple</b>
Price	\$24,000	<b>\$24,000 est</b>	\$3,200	<b>\$3,200 est</b>	\$29,000	<b>\$29,000 est</b>
Shelves compatible?	Yes	<b>Yes</b>	Yes	<b>Yes</b>	Yes	<b>Yes</b>
Previous generations read	1	<b>TBD</b>	2	<b>2</b>	1	<b>TBD</b>
Previous generations written	0	<b>TBD</b>	1	<b>1</b>	1	<b>TBD</b>
Bit Error Rate (BER)	1x10 <sup>-19</sup>	<b>1x10<sup>-19</sup></b>	<b>1x10<sup>-17</sup></b>	<b>1x10<sup>-17</sup></b>	1x10 <sup>-17</sup>	<b>1x10<sup>-17</sup></b>
Drive manufacturers	1	<b>1</b>	4+	<b>4+</b>	1	<b>1</b>
Availability	2008	<b>Late 2010</b>	2010	<b>2012</b>	2008	<b>2011</b>

## Oracle T10000B

The T10000B is the Oracle flagship high-capacity enterprise drive typically used in conjunction with Oracle robotic libraries, such as the SL8500. The EROS Center has eight T10000B drives for use in the SL8500, in addition to six first-generation T10000 drives (fig. 1).

### Advantages

- The T10000B is an evolution of the 9940, which the USGS has determined to be extremely reliable.
- Native capacity is 1 TB and native transfer rate is 120 MB/sec. The T10000B also can stream at 50 MB/sec, which is important because some disks may not be able to keep up at 120 MB/sec.
- The T10000B uses 32 channels per pass (compared to 16 on competing drives), which reduces serpentine passes. With 1,152 tracks, only 36 passes are required to read or write the entire tape.
- The T10000B is targeted to the enterprise storage market where data viability, speed, and capacity are more important than cost.
- The T10000B 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.
- T10000B drives provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- T10000B uses the same media as the T10000, allowing media re-use. Tapes written in T10000B format cannot be read by the T10000 drives.
- The T10000B has a 256 MB buffer, which prevents occasional data starvation from reducing the transfer rate.
- The Bit Error Rate (BER) is an industry best at  $1 \times 10^{-19}$ .
- A hardware encryption option module is available.

### Disadvantages

- The only cartridges available are produced for Oracle by Imation and Fuji.
- The T10000B drives are 7 times the price of the LTO5 but cheaper than the TS1130.
- Based on sales of the T10000 the T10000B sales are anticipated to be primarily for use in Oracle robotics. For this reason, the T10000B is anticipated to have a market share that will remain low compared to LTO, ensuring that media costs will remain high.
- The T10000B drive is only available from Oracle. This availability keeps the price high but does eliminate concerns of incompatibility.

### Summary

The T10000B is a high-capacity, high-transfer rate, enterprise-class drive for use in Oracle robotic libraries. The cost of media and drives far exceeds the cost of LTO, but media reuse for future generations 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, use of an enterprise solution for the second copy is not warranted.

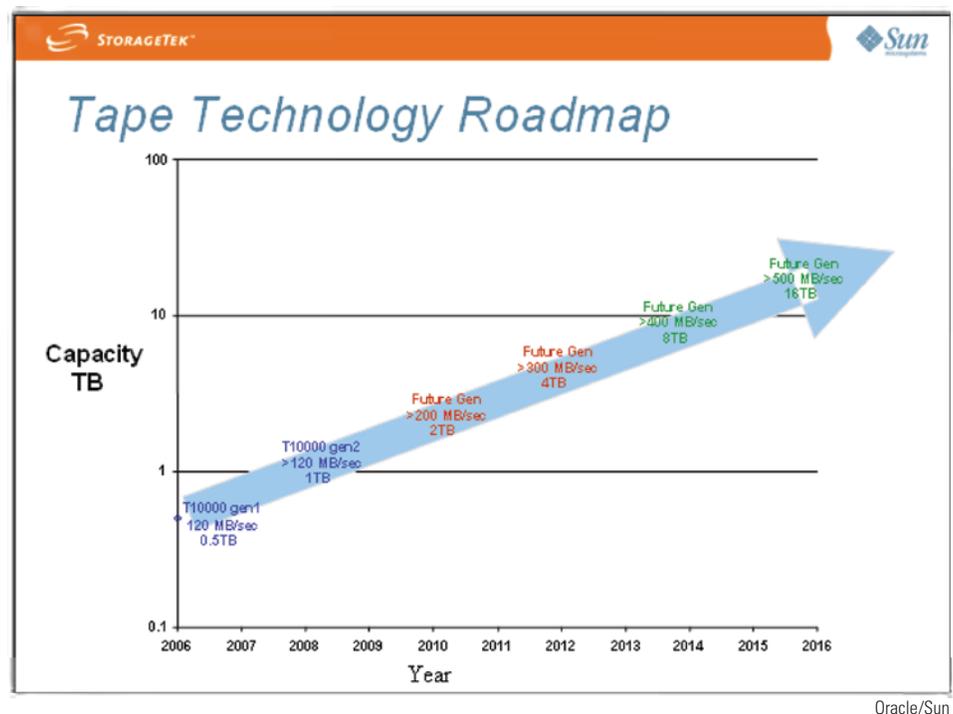


Figure 1. Screen capture showing the Oracle roadmap (uncompressed)

## Oracle T10000C

The T10000C is the third generation of the T10000 line. The T10000C was originally anticipated to ship in spring 2010, but is now expected to ship in late 2010 or early 2011.

### Advantages

- The T10000C is an evolution of the T10000/T10000B, which the USGS has determined to be extremely reliable.
- Native capacity is anticipated to be at least 2 TB and native transfer rate of at least 200+ MB/sec. The T10000C is expected to stream at lower rates, which is important because some disks may not be able to keep up at 200+ MB/sec.
- The T10000C is expected to use at least 32 channels per pass (compared to 16 on competing drives), which reduces serpentine passes.
- The T10000C is targeted to the enterprise storage market where data viability, speed, and capacity are more important than cost.
- The media for the T10000C is expected to differ from the media for the T10000/T10000B, but the T10000C may be able to read media written on T10000/T10000B. Like the T10000 media, the T10000C media will likely be 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.
- T10000C drives should provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- Some future follow-on drives are expected to use the same media, allowing media reuse.
- The T10000C is expected to have at least a 256 MB buffer, which prevents occasional data starvation from reducing the transfer rate.
- The BER is expected to be an industry best at  $1 \times 10^{-19}$ .
- A hardware encryption option module is anticipated.

### Disadvantages

- Cartridges may be supplied only by Oracle.
- The T10000C drives are expected to be 7 times the price of the LTO and cheaper than the TS1130.
- Based on sales of the T10000, the T10000C 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 T10000C drive is expected to be available only from Oracle. This availability keeps the price high but does eliminate concerns of incompatibility.

### Summary

The T10000C should replace the T10000/T10000B drives as the flagship high-capacity enterprise drive typically used in conjunction with Oracle robotic libraries because the T10000C should be priced comparably. The T10000C is not yet available and, therefore, was not assessed in the final evaluation.

### HP LTO5

The LTO5 is the most recent available generation of the LTO tape family (fig. 2) and will be tested at the EROS Center in late 2010.

### Advantages

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of 2006, LTO held an 82 percent market share. Since then, further development of the leading competing products (DLT and SAIT) has been discontinued. Since 2006, LTO drove DLT and SAIT-2 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.

### 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 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 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 observed two LTO1 incompatibility problems between HP and IBM: tapes written to end-of-tape (EOT) on the IBM cannot be read on the HP, and tapes written on the HP read at less than half speed on the IBM. Inter-brand incompatibilities can be avoided by using a single brand of drive.

## Summary

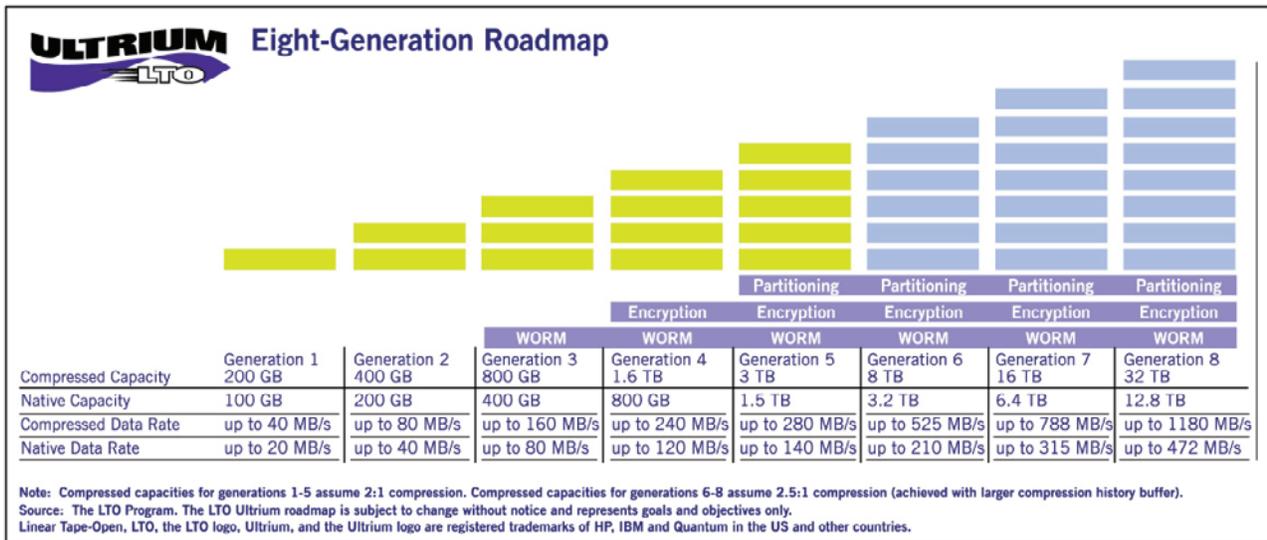
Testing of LTO5 technology at EROS will begin in late 2010. LTO has been reliable at USGS, with only a small number of failures commensurate with the design specifications for a mid-range tape technology.

## HP LTO6

The LTO6 is the next anticipated generation of the LTO tape family, with release anticipated in 2012 based on a typical LTO release cycle of 2 years.

## Advantages

- LTO has had phenomenal growth from the day of release in 2000; as of 2006, LTO held an 82 percent market share.
- 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.



LTO Consortium

Figure 2. Screen capture showing the LTO roadmap (with 2:1 compression)

### 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 usage is light.
- Each generation of LTO requires new media in order 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. EROS observed two LTO1 incompatibility problems between HP and IBM: tapes written to EOT on the IBM cannot be read on the HP, and tapes written on the HP read at less than half speed on the IBM. EROS resolved this issue by only deploying HP drives for production use.

### Summary

LTO6 is expected to be announced in 2011 and made available in 2012. LTO6 is not yet available and was not assessed in the final evaluation.

### IBM TS1130

The TS1130 is an enterprise-class tape drive, used primarily in robotic libraries and autoloaders. The TS1130 is a follow-on drive to the TS1120 (fig. 3).

### Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, and TS1120.
- Supports a 4 gigabit per second (Gbit/sec) Fiber Channel interface.
- Native capacity is 1 TB and native transfer rate is 160 MB/sec.
- The TS1130 is a robust storage technology, with the tape cartridge and drive built to withstand constant or frequent use in a robotic environment.
- The TS1130 uses the same media as the TS1120 and 3592, plus a new higher capacity cartridge.
- A hardware encryption feature is included in the drive.

### Disadvantages

- Designed primarily for use in IBM robotic libraries.

### Summary

The TS1130 does 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 USGS robotic library. IBM recently reported development of a recording method that will yield a capacity of 35 TB per cartridge, but IBM did not reveal a timeline.

		Gen 1	Gen 2	Gen 3	Gen 4	Gen 5	Gen 6
<b>3592 Model</b>		3592	TS1120	TS1130			
<b>M/T Model</b>							
<b>Native capacity</b>		300 GB	500 GB 700 GB	640 GB 1 TB	?	4 TB	8 TB
<b>Data transfer rate MB/S</b>		40	100	150 min	240 min	360 min	540 min
<b>With Compression</b>		Up to 120	Up to 300	380 +			
<b>Cartridge Type</b>		JJ/JA	JJ/JA	JJ/JA/JB	JA/JB/JC	JB/JC	JB/JC/JD
<b>WORM</b>		JR/JW	JR/JW	JR/JW/JX	JW/JX/JY	JX/JY	JX/JY/JZ
<b>Encryption</b>		N/A	Yes	Yes	Yes	Yes	Yes
<b>Sever Attachment</b>		Fibre FICON	Fibre FICON	Fibre FICON	Fibre FICON	Fibre FICON	Fibre FICON
		ESCON	ESCON				

Technology Demonstration  
1 TB - April, 2002  
8 TB - May, 2006

21 Statements of IBM future plans and directions are provided for information purposes only. Plans and direction are subject to change without notice. 3/4/2008 IBM Systems International Business Machines

Figure 3. Screen capture showing the IBM roadmap (uncompressed)

## IBM TS1140

The TS1140 is anticipated to be the next generation of the 3592 tape family, with release expected in 2011. Note that the TS1140 name has not been confirmed, but follows logically.

### Advantages

- Lineage includes the reliable 3480, 3490, 3590, 3592, TS1120, and TS1130.
- Should support a 4 or 8 Gbit/sec Fiber Channel interface.
- Native capacity is expected to be 2 TB and native transfer rate may exceed 240 MB/sec.
- The TS1140 will be a robust storage technology, with the tape cartridge and drive built to withstand constant or frequent use in a robotic environment.
- The TS1140 may use the same media as the TS1130.
- A hardware encryption feature should be included in the drive.

### Disadvantages

- Designed primarily for use in IBM robotic libraries.

### Summary

The TS1140 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 USGS robotic library. TS1140 is not yet available and was not assessed in the final evaluation.

## Tables

### Design Criteria

The design criteria and target market of a drive are inter-related (table 4). LTO5 is targeted to the backup market, as demonstrated by LTO marketing. The T10000B and TS1130 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 a secondary factor. Cost is not a major 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:

$$\begin{aligned} & ((100\text{-serpentine passes})/10)^+ \\ & (\text{absolute value of error rate exponent}/2)^+ \\ & (\text{construction } 3=\text{moderate usage, } 5=\text{high usage})^+ \\ & (\text{head contact } 3=\text{contact, } 5=\text{min contact}) \\ & / 2.59 \text{ (to adjust the highest rank to } 10) \end{aligned}$$

### Transfer Rate

Transfer rate is important because it establishes how quickly the migration of an archive dataset may be completed

**Table 4.** Design criteria and target market.

[Uncorrected error rates for some drives are not available but are presumed to be either the same as their predecessor or at least  $1 \times 10^{-17}$ . Yellow highlighted text indicates unverified information. MP, metal particle; TBD, to be determined; HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Technology	Serpentine tracks/passes	Target market	Tape composition	Uncorrected error rate	Cartridge construction rating	Head contact	Ranking
Oracle T10000B	1,152/36	Enterprise	Advanced MP	$1 \times 10^{-19}$	High usage	Minimum contact	10.0
Oracle T10000C	TBD	Enterprise	Advanced MP	$1 \times 10^{-19}$	High usage	Minimum contact	10.0
HP LTO5	1,280/80	Backup	Thin film MP	$1 \times 10^{-17}$	Moderate usage	Contact	6.4
HP LTO6	TBD	Backup	Thin film MP	$1 \times 10^{-17}$	Moderate usage	Contact	6.4
IBM TS1130	1,152/72	Enterprise	Advanced MP	$1 \times 10^{-17}$	High usage	Contact	8.2
IBM TS1140	TBD	Enterprise	Advanced MP	$1 \times 10^{-17}$	High usage	Contact	8.2

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 lack 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 half of the total read/write transfer rate of the leader would be ranked 5.

### 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 LTO4 at 757 GB of usable capacity per tape. The new 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 are not compressible, all capacities are native (uncompressed). Where measured capacities were not available, approximate capacities are determined based on the accuracy of specified capacities of previous generations.

The capacities listed in table 6 presume that a gigabyte = 1,073,741,824 bytes. The ratings were determined by computing each as the 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. Note that capacity yield varies by media vendor.

### Cost Analysis

Table 7 shows the relative drive and media costs, drive warranty, 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 found on the Web or on government price lists.

**Table 5.** Transfer rates.

[Yellow highlighted text indicates unverified information. EROS, Earth Resources Observation and Science; HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines; MB/sec, megabyte per second]

Tape drive technology	Advertised/proposed native rate	Source of test results	Actual/approximate native write transfer rate	Advertised capacity (in percent)	Actual/approximate native read transfer rate	Advertised capacity (in percent)	Ranking
Oracle T10000B	120 MB/sec	EROS testing	109.00 MB/sec	91	120.00 MB/sec	100	7.4
Oracle T10000C	180 MB/sec	Approximate	163.80 MB/sec	91	180.00 MB/sec	100	7.4
HP LTO5	140 MB/sec	Approximate	126.70 MB/sec	90.5	126.56 MB/sec	90.4	8.2
HP LTO6	210 MB/sec	Approximate	190.05 MB/sec	90.5	189.84 MB/sec	90.4	8.2
IBM TS1130	160 MB/sec	Vendor	153.92 MB/sec	96.2	153.92 MB/sec	96.2	10.0
IBM TS1140	240 MB/sec	Approximate	230.88 MB/sec	96.2	230.88 MB/sec	96.2	10.0

**Table 6.** Storage capacities.

[Yellow highlighted text indicates unverified information. TB, terabyte; GB, gigabyte, HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Tape drive technology	Advertised/proposed native capacity	Measured/approximate native capacity	Advertised capacity (in percent)	Ranking
Oracle T10000B	1.0 TB	936 GB	93.6	6.6
Oracle T10000C	2.0 TB	1,872 GB	93.6	6.6
HP LTO5	1.5 TB	1,420 GB approximate	94.7 approximate	10.0
HP LTO6	3.2 TB	3,030 GB approximate	94.7 approximate	10.0
IBM TS1130	1.0 TB	950 GB approximate	95.0 approximate	6.7
IBM TS1140	2.0 TB	1,900 GB approximate	95.0 approximate	6.7

**Table 7.** Drive and media costs.

[Yellow highlighted text indicates unverified information. TB, terabyte; est, estimated; HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Technology	Drive (dollars per each)	Drive warranty	Media (dollars per each)	Media (dollars per TB)	Ranking drive cost	Ranking media cost per TB
Oracle T10000B	\$24,000	12 month	\$125	\$125	1.3	5.6
Oracle T10000C	\$24,000 est	12 month	\$125 est	\$62 est	1.3	5.6
HP LTO5	\$3,200	36 month	\$105	\$70	10.0	10.0
HP LTO6	\$3,200 est	36 month	\$105 est	\$33 est	10.0	10.0
IBM TS1130	\$29,000	12 month	\$178	\$178	1.1	3.9
IBM TS1140	\$29,000 est	12 month	\$178 est	\$89 est	1.1	3.9

Maintenance should be a consideration but was removed from this iteration of the study because of the tenuous status of Oracle support costs as of this writing, and incomplete information on LTO support. Maintenance will be reconsidered in the next update.

### 3.5 Scenarios

Table 8 shows the total drive and media cost for three scenarios. These scenarios presume that each dataset or project stands alone, although pooling resources for multiple datasets can mitigate cost. Competition often results in a considerable drop in media prices within 6 months after product introduction.

Rankings are based on the 100TB option and were established by setting the cheapest to 10 then rating each of the others against the lowest cost. Advertised/proposed 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).

[Yellow highlighted text indicates unverified information. TB, terabyte; HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Technology	100 TB 2 drives	200 TB 3 drives	400 TB 4 drives	100 TB ranking
Oracle T10000B	\$60,500	\$97,000	\$146,000	2.2
Oracle T10000C	\$54,250	\$84,500	\$121,000	2.2
HP LTO5	\$13,435	\$23,565	\$40,730	10.0
HP LTO6	\$9,700	\$16,200	\$26,000	10.0
IBM TS1130	\$75,800	\$122,600	\$187,200	1.8
IBM TS1140	\$66,900	\$104,800	\$151,600	1.8

### Vendor Analyses

Table 9 provides an analysis of each company and the stability of each technology. 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 T10000B, the technology was based on the predecessor 9940; therefore, the technology age includes the 9940. The longevity rankings were determined by the following formula:

$$\frac{\text{(company age + technology age)}}{11.4} \text{ (to adjust the highest rank to 10)}$$

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

**Table 9.** Vendor analyses.

[STK, StorageTek; HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Company	Technology	Years in business	Technology age, in years	Longevity ranking
Oracle/Sun/STK	T10000	41 (1969)	10 (2000)	4.5
HP	LTO	71 (1939)	10 (2000)	7.1
IBM	3592 (3590)	99 (1911)	15 (1995)	10.0

### Drive Compatibility

Table 10 shows the level of intergeneration drive compatibility and the future drives planned. The columns “Percentage of previous generations read” and “Percentage of 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:

$$\frac{(\text{Percentage of previous generations read} + \text{Percentage of previous generations written} + (\text{Future generations planned} \times 20))}{21} \text{ (to adjust the highest rank to 10)}$$

**Table 10.** Drive compatibility.

[Yellow highlighted text indicates unverified information. HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Technology	Percentage of previous generations read	Percentage of previous generations written	Future generations planned	Ranking
Oracle T10000B	100	0	3	7.6
Oracle T10000C	100	0	2	7.6
HP LTO5	50	25	3	6.4
HP LTO6	40	20	2	6.4
IBM TS1130	100	50	3	10.0
IBM TS1140	100	50	2	10.0

### Ranking Summary

The ranking summary provides a quick reference to the rankings (table 11).

## Conclusions and Recommendations for USGS Offline Archiving Requirements

### Weighted Decision Matrix

Table 12 provides a weighted analysis of the drives considered. 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 in green are relative ratings for each technology. The columns in yellow 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 over 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).

**Table 11.** Ranking summaries.

[Blue indicates the highest ranking in category. HP, Hewlett-Packard; LTO, Linear Tape-Open; IBM, International Business Machines]

Drive	Design criteria	Capacity	Media cost	Drive compatibility	Transfer rate	Drive cost	Vendor analyses	Scenario cost
T10000B	10.0	6.6	5.6	7.6	7.4	1.3	4.5	2.2
HP LTO5	6.4	10.0	10.0	6.4	8.2	10.0	7.1	10.0
IBM TS1130	8.2	6.7	3.9	10.0	10.0	1.1	10.0	1.8

**Table 12.** Weighted decision matrix.

[TB, per terabyte]

Selection criteria	Weight	Oracle T10000B	HP LTO5	IBM TS1130	Oracle T10000B	HP LTO5	IBM TS1130
Design criteria	0	10.0	6.4	8.2	0.0	0.0	0.0
Capacity	20	6.6	10.0	6.7	132.0	200.0	134.0
Media cost per TB	0	5.6	10.0	3.9	0.0	0.0	0.0
Compatibility	15	7.6	6.4	10.0	114.0	96.0	150.0
Transfer rate	15	7.4	8.2	10.0	111.0	123.0	150.0
Drive cost	0	1.3	10.0	1.1	0.0	0.0	0.0
Vendor analyses	15	4.5	7.1	10.0	67.5	106.5	150.0
Scenario cost	35	2.2	10.0	1.8	77.0	350.0	63.0
Total weighted score					501.5	875.5	647.0

- 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 measured or approximate capacity is used rather than advertised capacity (table 8).

Note that 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.

## Conclusions and Notes

LTO5 achieved the highest total score in the study; therefore, no compelling reason exists to abandon LTO to adopt a new standard offline archive technology.

LTO5 and TS1130 were not available to be tested for this study; therefore, performance and capacity figures were based on vendor or customer 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 T10000B, as are most archives at USGS.
- As any drive saturates the market, media and drive costs drop. Based on USGS experience with enterprise tape technology and observation of Oracle and IBM

pricing, enterprise drives such as the T10000B and TS1130 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.

## Recommendations

1. The USGS should continue with LTO4 as the offline storage media of choice, then test and move to LTO5, when available.
2. Data stored on LTO2 and LTO3 should be migrated to LTO5 in the next 2 years.
3. To reduce risk, the USGS should continue the strategy of storing datasets on multiple technologies when only two copies exist. For example, store a working copy of a dataset on nearline T10000B and offline/offsite copies on LTO. 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.
5. The USGS should adopt a policy of periodically testing archive tapes for readability. 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 and the checksum stored in the corresponding inventory record. When a file is retrieved from either the Silo or the offline media, integrity can be verified. Verification of each retrieved file may not be feasible because of CPU impacts.
7. All data should be migrated to new media from 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. This is a best practice supported by NARA.
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. 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 a benefit, even if the technology is not the leading technology continuously. 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.
10. The USGS should plan to update this trade study periodically. Annually may be too frequent to observe market changes because drives are typically updated on a 2- or 3-year cycle.



# Appendix: Citations

---

## Vendor Sites

<http://h18006.www1.hp.com/storage/tapestorage/tapedrives.html> (Hewlett-Packard)

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

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

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

<http://www.informationweek.com/news/showArticle.jhtml?articleID=187203674>

<http://www.techworld.com/storage/news/index.cfm?newsID=5888&pagetype=samechan>

[http://en.wikipedia.org/wiki/Holographic\\_Versatile\\_Disc](http://en.wikipedia.org/wiki/Holographic_Versatile_Disc)

<http://www.norsam.com/hdrosetta.htm>

[http://www.theregister.co.uk/2007/03/21/lto\\_beats\\_dlt/](http://www.theregister.co.uk/2007/03/21/lto_beats_dlt/)

<https://www.bluestoragemedia.com/External/BigBlueBytes/Product%20Information/3592%20Gen%202/IBM%20System%20Storage%20TS1120%20Tape%20Drive%20Training%20Presentation.pdf>

<http://www.techworld.com/storage/features/index.cfm?featureid=3728>

[http://www.imation.com/euc/pdfs/EUC\\_07\\_Kenyon.pdf](http://www.imation.com/euc/pdfs/EUC_07_Kenyon.pdf)

<http://www-03.ibm.com/systems/storage/tape/ts1130/index.html>

<http://h10010.www1.hp.com/wwpc/us/en/sm/WF25a/12169-304612-3446236-3446236-3446236-4150338.html>

<http://dlc.sun.com/pdf/316194802B/316194802B.pdf>

**18 Archive and Records Management—Fiscal Year 2010 Offline Archive Media Trade Study**

*<http://www.lto-technology.com/pdf/2006-7-25.pdf>*

*<http://itknowledgeexchange.techtarget.com/storage-soup/hp-quantum-in-cahoots-for-lto-5/>*

*<http://www.itjungle.com/fhs/fhs020508-story10.html>*

*[http://searchdatabackup.techtarget.com/news/article/0,289142,sid187\\_gci1355225,00.html](http://searchdatabackup.techtarget.com/news/article/0,289142,sid187_gci1355225,00.html)*

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

*[http://en.wikipedia.org/wiki/Linear\\_Tape-Open](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>*

*<ftp://service.boulder.ibm.com/storage/tape/clipper.pdf>*

Publishing support provided by:  
Rolla and Lafayette Publishing Service Centers

For more information concerning this publication, contact:  
U.S. Geological Survey Earth Resources Observation and Science  
(EROS) Center  
47914 252nd Street  
Sioux Falls, South Dakota 57198  
(605) 594-6151

Or visit the EROS Center Web site at:  
World Wide Web: <http://eros.usgs.gov/>

