Technical Support Services Contract (TSSC)					
USGS EROS	**O8HQCN0005**	Deliverable ID	TSSC-00-02-4.2-0001		
FY08 Archive Media Trade Study		Date	06/25/08	Version Nbr	1.0

ARCHIVE AND RECORDS MANAGEMENT

FY08 ARCHIVE MEDIA TRADE STUDY

STINGER GHAFFARIAN TECHNOLOGIES

ADNET SYSTEMS, INC. Science and Information Technology

NCDC IMAGING

Riverside Technology, inc. Integrated Solutions in Science, Engineering and Information Technology

INFORMATION DYNAMICS

Offline Archive Media Trade Study

June 2008

By SGT

Prepared by:

Tom Bodoh Date
Sr. Systems Engineer
SGT

Reviewed by:

Michael North Date
Work Manager
SGT

Reviewed by:

Ken Gacke Date
Sr. Systems Engineer
SGT

Reviewed by:

Cheryl Greenhagen Date
Sr. Systems Engineer
SGT

Reviewed by:

Al Engelbrecht Date
Principal Electronic Engineer
SGT

Reviewed by:

Kenneth Boettcher Date
Work Manager, Long Term Archive
SGT

Preface

This document contains the Offline Archive Media Trade Study prepared by SGT for the USGS. The Trade Study presents the background, technical assessment, test results, and recommendations.

The USGS uses trade studies and reviews for internal purposes and does not endorse vendors or products.

This Trade Study was prepared by

Stinger Ghaffarian Technologies (SGT) under contract 08HQCN0005
USGS/EROS
Sioux Falls, SD 57198

Team members:

Tom Bodoh
Ken Boettcher
Ken Gacke
Cheryl Greenhagen
Al Engelbrecht

Abstract

This document is a trade study comparing offline digital archive storage technologies. The document compares and assesses several technologies and recommends which could be deployed as the next generation standard for the U.S. Geological Survey (USGS) Earth Resources Observation and Science (EROS) Center. Archives must regularly migrate to the next generation of digital archive technology, and the technology chosen must maintain data integrity until the next migration. This document is the FY08 revision of a study completed in FY01 (Fiscal Year 2001) and revised in FY03, FY04, and FY06.

Revision History

February 2004

- Added revision history page
- 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 SuperDLT 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 specs 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 SAIT2 as future technologies
- Made vendor analyses formula more equitable, increasing weight 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

Contents

1.0 Introduction

1.1 Purpose and Scope

Typically, the purpose of a trade study is to analyze different courses of action and to provide the necessary information for the sponsor to reach a conclusion. In other cases, a trade study revalidates 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 migration. Data should be migrated before reliability 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. There continues to be a need for offline storage for infrequently used working copies, as well as master and offsite copies where the working copy is stored nearline.

Note that LTO4 is the current archive media of choice at EROS. There is no compelling reason for the USGS to change technologies at this time, and given the advantages of intergeneration 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. As with all technologies, eventually LTO will no longer meet EROS requirements, and this study (in future revisions) will have shown the way to the obvious emerging replacement.

1.2 Background

The USGS Earth Resources Observation and Science (EROS) Center, located in Sioux Falls, SD, has archived offline datasets using several technologies.

Tape Drive Technology	Capacity	Transfer rate	Type
HDT	3.4 GB	10.6 MB/sec	Analog
3480	200 MB	2.0 MB/sec	Digital
3490	900 MB	2.7 MB/sec	Digital
DLT 7000	32 GB	4.7 MB/sec	Digital
DCT (Ampex DCRsI)	45 GB	12.0 MB/sec	Analog
SuperDLT 220	98.8 GB	8.1 MB/sec	Digital
HP LTO Ultrium 2	**197 GB**	**23.8 MB/sec**	**Digital**
HP LTO Ultrium 3	**378 GB**	**31.0 MB/sec**	**Digital**
HP LTO Ultrium 4	**757 GB**	**108.0 MB/sec**	**Digital**

Table 1-1 Recent and Current Archive Technologies Used at EROS

(Current in bold)

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. This migration freed up enough library shelving to ensure that the library should never need to be expanded, and may in fact be reduced in size. The USGS uses LTO3 extensively for onsite backups. HDT, 3480/3490, and 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.

1.3 Data reliability

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

- The number of archival copies: The dependence on the master copy and the level of risk rise when a working copy is not robust. The master and working copies need not be on similar media. USGS archives typically have both working and master copies, and an offsite copy is desirable. A slightly less reliable drive technology can be used if there are a sufficient number of copies of the archive or if one copy uses an enterprise or archive drive technology.

- The storage location and environment: This is a constant for all of the technologies assessed since all media is stored in a secure and climate-controlled environment.

- The composition of the media: Some media compositions last much longer than others, though all of 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 over an error. Though error detection upon write is required, additional attention to data recovery upon read is a higher priority since media degradation will lead to eventual read errors.

- Primary Market: This criterion describes the target market of a drive and the characteristics of drives within that market.

    - A drive targeted to the backup market is designed for write many/read rarely and depends more on write error detection since the data is 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 upon 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 many, and more emphasis would be placed on detecting and correcting errors upon read—though there are currently no drives designed or marketed primarily for archive. Most vendors would argue that their products are archive devices, but if forced to choose their primary market no vendor would choose the very limited archive market over the lucrative backup market.

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	Up to 100% duty cycle for drives and media used with robotics
Archive	High	Write once, read many	Long-term reliability

Table 1-2 Tape Drive Markets and Characteristics

The reliability of a long-term archive technology relates primarily to the long-term viability of the recorded media. Since it is wise to implement a technology early enough in its life cycle that drives can be kept viable through the lifetime of a given media (or replaced with newer backward-compatible models), a definitive leader in reliability is difficult to determine except in retrospect. 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 Sun/StorageTek products to be very reliable, though the Sun/StorageTek D3 helical scan drive was problematic and was discontinued quickly. On two occasions, 9840 tapes that encountered unrecoverable errors were sent to StorageTek for recovery. One tape was recovered, but the other was unrecoverable due to cartridge contamination. The LP DAAC has experienced problems with replacement 9940B drives coming from the Philippines plant. One T10000 tape recently failed to read, and upon analysis it was found that the imbedded Radio-frequency Identification (RFID) chip had failed. Sun was able to recover the data.

1.4 Selection criteria

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

1. The technology must be currently available and shipping to be considered in the final analysis. It also must be the latest drive in the line. Drives that are anticipated/announced but not shipping are mentioned but not ranked in the final analysis.

2. The technology must hold at least 500 GB of uncompressed data.

3. The technology must have an uncompressed write transfer rate of at least 60 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 since it 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. Allowing only three to five years between migrations would better ensure the availability of drives and decrease the likelihood of media deterioration.

5. The technology must not be hampered by a poor reliability or performance history. Helical scan technologies such as 8mm, 4mm, DAT, and D3 have proven unreliable in the past.

The following currently available drive technologies were selected for consideration.

1. Sun T10000

2. HP LTO4 (Linear Tape Open)

3. Quantum DLT-S4 (Digital Linear Tape)

4. International Business Machines (IBM) TS1120

1.5 Dismissed technologies

The following technologies were dismissed from further analysis or consideration.

1.5.1 Disk

Disk prices continue to drop, while reliability, performance, and capacity increase. Cost, management overhead, cooling, and power are barriers to using disk to archive large datasets. In addition, there is inherent risk that an archive dataset stored online without an offline copy may be lost due to intentional or unintentional corruption. While tape could be kept viable up to ten years, the more costly disk is typically replaced every four or five years in order to maintain supportability, reliability, and performance. There is benefit in serving frequently used working copies on disk, though there must be an offline master copy.

1.5.2 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 based on obvious limitations in performance, capacity, or reliability. These products are not a good fit for large digital archives.

1.5.3 Sun (STK) 9840

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

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

Exabyte has evolved its 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. While media costs are low, transfer rates are low and the USGS experience with consumer-grade storage technologies has shown that they cannot withstand the rigors of a long-term archive.

Tape drives such as the 8mm/Exabyte, which became popular in the 1990s, were based on consumer-grade helical scan technology and were notably slow and unreliable. Long start/stop times dictated that if data was not kept streaming, then the effective transfer rate dropped drastically. The necessarily complex drive path led to problems with 8mm drives mangling tapes and a confusing array of firmware versions often yielded unpredictable behavior and hangs. The STK foray into helical scan was short lived due to irreconcilable problems. The transition from a market once ruled by 4mm/8mm helical scan drives to one ruled by LTO/DLT occurred 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 relatively 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, HD-DVD, Blu-Ray

Digital Video Disc (DVD) and related technologies seem promising from the standpoint of longevity of the media, although 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.

HD-DVD recently 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 CD and DVD, Blu-Ray suffers from high media costs and low transfer rates, and given optical media history, its 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, to date none are shipping products.

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 currently indicate a capacity up to 3.9 TB per disc and a transfer rate of 125 MB/sec. Rivals claim up to 100 TB per disc will be possible. Like other optical products, the cost model will likely be prohibitive compared to tape.

2.0 Technical Assessment

2.1 Analysis

This technical assessment includes drives selected for final evaluation (T10000, LTO4, DLT-S4, TS1120), as well as drives anticipated to be released in the near future (T10000B, LTO5, TS1130). 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 assessed, though only the bolded drives will be included in the final evaluation:

- **Sun T10000**
- Sun T10000B
- **HP LTO4**
- HP LTO5
- **Quantum DLT-S4**
- **IBM TS1120**
- IBM TS1130

Specification	T10000	T10000B	HP LTO4	HP LTO5	DLT-S4	TS1120	TS1130
Uncompressed capacity	500 GB	1.0 TB	800 GB	1.6 TB	800 GB	700 GB	1.0 TB
Uncompressed xfer rate	120 MB/sec	120 MB/sec	120 MB/sec	180 MB/sec	60 MB/sec	104 MB/sec	150 MB/sec
Recording technology	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine	Serpentine
Tracks	768	1152	896	TBD	1280	896	TBD
Channels	32	32	16	TBD	16	16	TBD
Passes	24	36	56	TBD	80	56	TBD
Tape velocity	4.95 m/sec	3.74 m/sec	6.2 m/sec	TBD	TBD	6.2 m/sec	TBD
Type	Enterprise	Enterprise	Backup	Backup	Backup	Enterprise	Enterprise
Encryption support	HW option	HW option	HW option	TBD	Software	HW built-in	HW built-in
Buffer size	256 MB	256 MB	128 MB*	256 MB	256 MB	512 MB	512 MB
Adaptive speeds	2 speeds	2 speeds	40 to 120	TBD	None	6 speeds	TBD
Price	$24k	$24k est	$2.5k	$4k est	$1.6k	$26k	$26k est
Shelves compatible?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prev generations read	NA	1	2	2	2	1	TBD
Prev generations written	NA	0	1	1	0	1	TBD
Bit Error Rate (BER)	$1x10^{-19}$	$1x10^{-19}$	$1x10^{-17}$	$1x10^{-17}$	$1x10^{-17}$	$1x10^{-17}$	$1x10^{-17}$
Drive manufacturers	1	1	4	4	1	1	1
Availability	Now	Oct 2008	Now	Dec 2009	Now	Now	July 2008

Table 2-1 Technology comparison

(yellow highlighted text indicates unverified information)

*** Tandberg and IBM LTO4 drives have a 256 MB buffer**

Sun T10000:

The T10000 is Sun's flagship high-capacity enterprise drive typically used in conjunction with Sun robotic libraries, such as the SL8500. EROS has seven T10000 drives for use in the SL8500.

Advantages:

- The T10000 is an evolution of the 9940, which the USGS has found to be to be extremely reliable.
- Native capacity is 500 GB and native transfer rate is 120 MB/sec. It can also stream at 50 MB/sec, which is important since some disks may not be able to keep up at 120 MB/sec.
- The T10000 uses 32 channels per pass (vs. 16 on competing drives), which reduces serpentine passes. With 768 tracks, only 24 passes are required to read or write the entire tape.
- The T10000 is targeted to the enterprise storage market where data viability, speed, and capacity are more important than cost.
- The T10000 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 due to their durability.
- T10000 drives provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- Sun indicates that the follow-on drive will use the same media, allowing media reuse.
- The T10000 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 $1x10^{-19}$.
- A hardware encryption option module is available.

Disadvantages:

- Though announced by Fuji, T10000 compatible media has never materialized. The only cartridges available are produced for Sun by Imation.
- While the T10000 is an evolution of the 9940, the media is incompatible. Consequently, 9940 media cannot be reused or read by the T10000.
- The T10000 drives are 10 times the price of the LTO4 but comparable to the TS1120.
- Based on sales of past Sun drives such as the 9840 and 9940, it is anticipated that sales of the T10000 will be primarily for use in Sun robotics. For this reason, it is anticipated that market share will remain low compared to LTO.
- The T10000 drive is only available from Sun. This keeps the price high but does eliminate concerns of incompatibility.

Summary:

The T10000 is a high-capacity, high-transfer rate enterprise-class drive for use in robotic libraries. The cost of media and drives far exceeds the cost of most competing products, though media reuse for future generations would effectively reduce media costs. The robust technology would be a prime choice 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 T10000, use of an enterprise solution for the second copy is not warranted. EROS has experienced one T10000 tape failure, when an RFID chip failed, rendering the tape unreadable. The data was recovered by Sun after installing a new RFID chip.

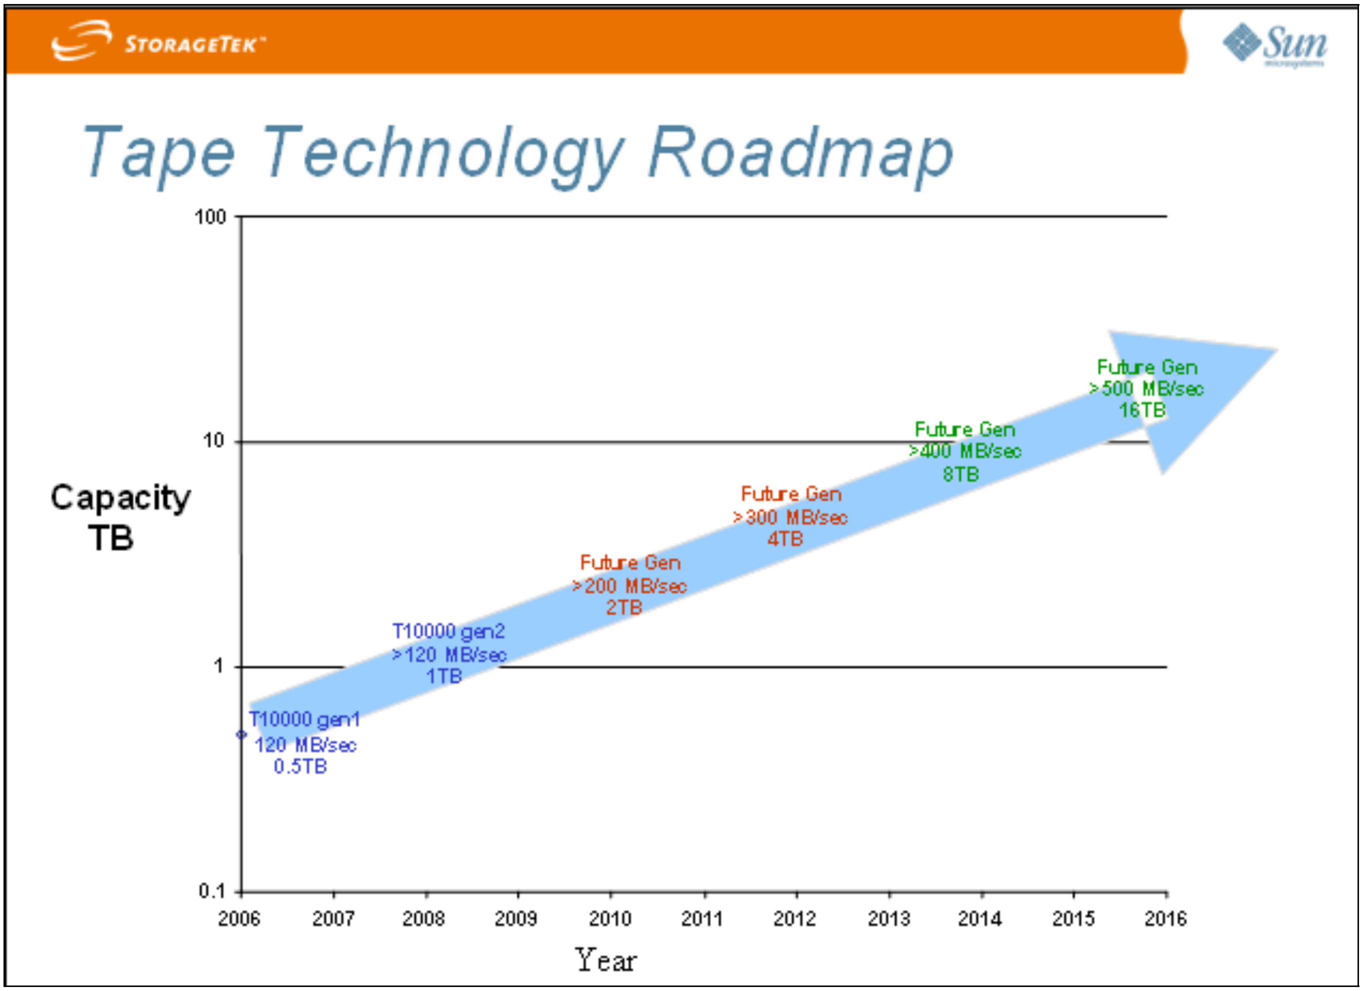

Figure 2-1 Sun Roadmap (uncompressed)

(source: Sun)

Sun T10000B:

The T10000B is the second generation of the T10000 line. It was originally announced to ship in May 2008, but problems were encountered in testing. It is now expected to ship in October, 2008.

Advantages:

- The T10000B is an evolution of the T10000, which the USGS has found to be to be extremely reliable.
- Native capacity is 1 TB and native transfer is at least 120 MB/sec. It is expected to stream at lower rates, which is important since some disks may not be able to keep up at 120 MB/sec.
- The T10000B uses 32 channels per pass (vs. 16 on competing drives), which reduces serpentine passes. With 1,152 tracks expected, 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 T10000 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 due to their durability.
- T10000B drives provide drive statistics for servo errors, bytes read/written, I/O retries, and permanent errors.
- Sun indicates that follow-on drives will use the same media, allowing media reuse.
- 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 $1x10^{-19}$.
- A hardware encryption option module is available.

Disadvantages:

- Though announced by Fuji, T10000 compatible media has never materialized. The only cartridges available are produced for Sun by Imation.
- The T10000B drives are expected to be 5–13 times the price of the SDLT or LTO but comparable to the TS1120.
- Based on sales of past Sun drives, it is anticipated that sales of the T10000B will be primarily for use in Sun robotics. For this reason, it is anticipated that market share will remain low compared to LTO.
- The T10000B drive is only available from Sun. This keeps the price high but does eliminate concerns of incompatibility.

Summary:

Upgrading to the T10000B would free up tapes and slots, and increase performance. The T10000B should replace the T10000 drive as the flagship high-capacity enterprise drive typically used in conjunction with Sun robotic libraries since it will use the same media and should be priced comparably. Migration from T10000 to T10000B will be made easier since the media is the same, the T10000B will read T10000 formatted tapes, and most HSM systems will support automated background migration. Given that the T10000B does not yet exist as a product, it cannot be one of the technologies assessed in the final evaluation.

HP LTO4:

The LTO4 is the current generation of the LTO tape family and is the current offline archive media of choice at EROS.

Advantages:

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of 2006, it held an 82% market share.
- Native capacity is 800 GB and native transfer rate is 120 MB/sec.
- The HP LTO4 drive has the capability to adapt the transfer rate to match the streaming speed of the source.
- LTO4 is backward read compatible with LTO2 and LTO3, and backward write compatible with LTO3.
- 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 LTO4 has a 128 MB or 256 MB buffer, which prevents occasional data starvation from reducing the transfer rate. The buffer size depends on the brand and model.
- A hardware encryption option is available.

Disadvantages:

- LTO is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write many/read rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.
- Repeated end-to-end use of a tape would be a concern since one end-to-end read/write incurs 56 passes (896 tracks divided by 16 channels).
- When STK adopted LTO technology for support by their Powderhorn robotic library, the robotic hand pressure was reduced to prevent crushing the thin LTO shell. With the decreased hand pressure, the speed of the robotic arm was reduced to prevent the arm from throwing tapes. This reduced the tape exchange rate of the Powderhorn.
- Each generation of LTO requires new media, ensuring that media costs will be significantly higher until market saturation drives the price down.
- LTO was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant use. EROS has experienced more LTO tape failures than with 9840/9940/T10000.
- LTO was co-developed by IBM, HP, and Quantum (acquired from Seagate/Certance). This type of partnership makes it possible for each vendor to interpret the specifications differently and to design drives which 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:

Adoption of LTO4 technology at EROS began in 2007. Three LTO4 drives now support the generation and ingest of archive media. LTO has been quite reliable at EROS, with only a small number of failures commensurate with the design specifications for a mid-range tape technology. During a recent migration from an old HSM to a new one, time was saved by re-ingesting data from LTO rather than performing a network transfer between the HSM systems.

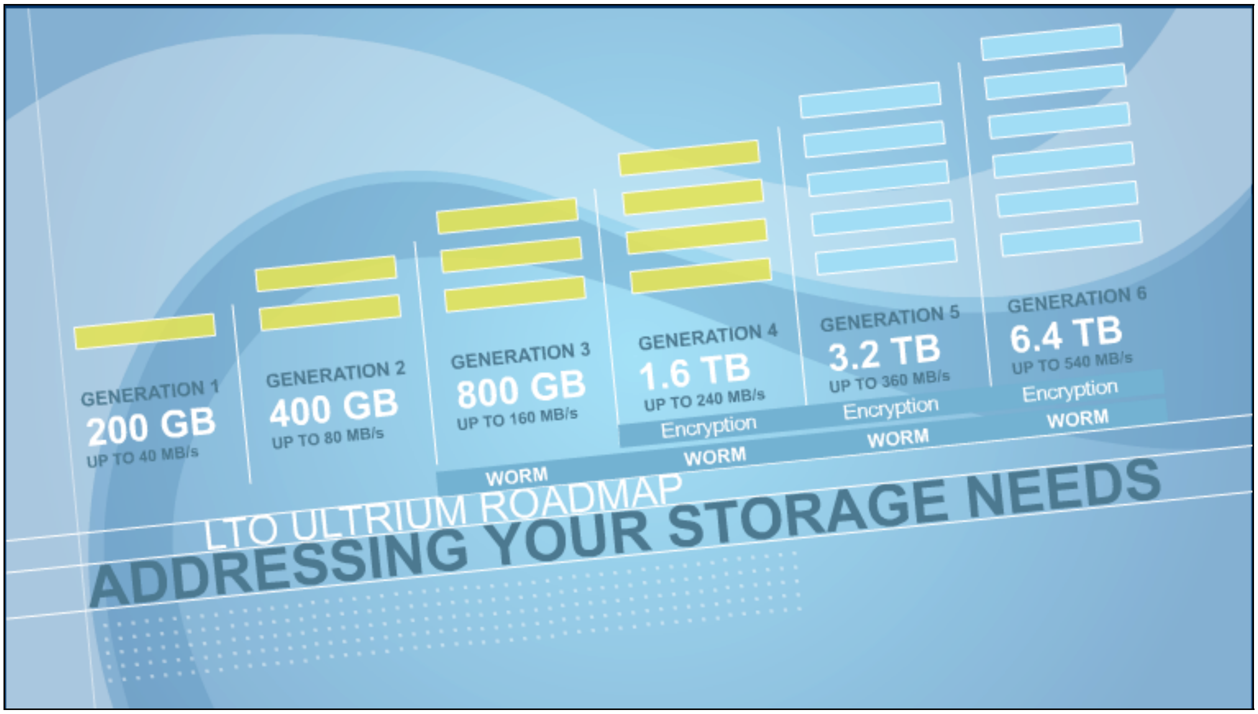

Figure 2-2 LTO Roadmap (with 2:1 compression)
(source: LTO Consortium)

HP LTO5:

The LTO5 is the next anticipated generation of the LTO tape family, with release anticipated in late 2009.

Advantages:

- LTO has enjoyed phenomenal growth from the day of release in 2000; as of 2006, it held an 82% market share.
- Native capacity is expected to be 1.6 TB and native transfer rate is expected to be 240 MB/sec.
- The HP LTO5 drive is anticipated to use an adaptive transfer rate to match the streaming speed of the source.
- LTO5 should be backward read compatible with LTO3 and LTO4, and backward write compatible with LTO4.
- 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.
- A hardware encryption option is anticipated.

Disadvantages:

- LTO is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write many/read rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.
- Repeated end-to-end use of a tape would be a concern since one end-to-end read/write is expected to incur 56 or more passes.
- When STK adopted LTO technology for support by their Powderhorn robotic library, the robotic hand pressure was reduced to prevent crushing the thin LTO shell. With the decreased hand pressure, the speed of the robotic arm was reduced to prevent the arm from throwing tapes. This reduced the tape exchange rate of the Powderhorn.
- Each generation of LTO requires new media, ensuring that media costs will be significantly higher until market saturation drives the price down.
- LTO was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant use. EROS has experienced more LTO tape failures than with 9840/9940/T10000.
- LTO was co-developed by IBM, HP, and Quantum (acquired from Seagate/Certance). This type of partnership makes it possible for each vendor to interpret the specifications differently and to design drives which 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:

It is anticipated that LTO5 will be announced in spring 2009 and made available by the end of 2009. In January 2008, the LTO consortium began offering product development LTO5 licenses. Given that the LTO5 does not yet exist as a product, it cannot be one of the technologies assessed in the final evaluation.

Quantum DLT-S4:

Several years ago, SDLT overtook 4mm and 8mm helical scan as the technology of choice for backups.

Advantages:

- Native capacity is stated to be 800 GB and native transfer rate is 60 MB/sec.
- The DLT-S4 is backward read compatible with SDLT 320 and SDLT 600.

Disadvantages:

- DLT-S4 uses different media than previous generations, ensuring that media costs will be significantly higher until market saturation drives the price down.
- DLT-S4 is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write many/read rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.
- Repeated end-to-end use of a tape would be a concern since each end-to-end use would incur 80 passes.
- DLT-S4 was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant use.
- The DLT-S4 is not write compatible with any previous generations.
- The DLTSage software suite includes software data encryption which is CPU intensive.

Summary:

In 2004, Quantum acquired Certance in order to enter the LTO marketplace as a hedge against the probable demise of DLT, given the market dominance of LTO. According to a 3/21/07 article on theregister.co.uk, Quantum's CEO admitted that because DLT had lost market dominance to LTO they may never release any new DLT technologies. Since there have been no further announcements, it would seem that DLT-S4 will be the last model in the DLT line, though the DLT-S4 may generate some residual sales until LTO5 appears. The link for the DLT roadmap on the Quantum Web site is now defunct.

IBM TS1120:

The TS1120 is an enterprise-class tape drive, used primarily in robotic libraries and autoloaders. It is a follow-on drive to the 3592.

Advantages:

- Lineage includes the very reliable 3480, 3490, 3590, and 3592.
- 4 Gbit/sec Fiber Channel interface.
- Native capacity is stated to be 500 GB and native transfer rate is 104 MB/sec.
- The TS1120 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 TS1120 uses the same media as the 3592.
- A hardware encryption feature is included in the drive.

Disadvantages:

- Designed primarily for use in IBM robotic libraries.

Summary:

The TS1120 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 EROS robotic library. IBM has announced a prototype tape technology which would be capable of storing 8 TB per cartridge and is planned to be available by 2011. IBM is expected to ship the 1 TB TS1130 in 2008.

IBM

3592/TS1120 Tape Drive Roadmap

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

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

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

Figure 2-3 IBM Roadmap (uncompressed)

(Source: IBM)

IBM TS1130:

The TS1130 is the next anticipated generation of the 3592 tape family, with release anticipated in mid 2008.

Advantages:

- Lineage includes the very reliable 3480, 3490, 3590, 3592, and TS1120.
- Should support a 4 Gbit/sec Fiber Channel interface.
- Native capacity is expected to be 1 TB and native transfer rate may exceed 150 MB/sec.
- The TS1130 will be 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 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. Given that the TS1130 does not yet exist as a product, it cannot be one of the technologies assessed in the final evaluation.

3.0 Tables

3.1 Design criteria

The design criteria and target market of a drive are interrelated. Drive technologies such as LTO and SDLT are targeted to the backup market, as demonstrated by their marketing. The T10000 and TS1120 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 since the data is 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 upon read and write. Enterprise drives are typically built for reliability and speed, with capacity a secondary factor. Cost is a 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 many, and more emphasis would be placed on detecting and correcting errors upon read; however, there are currently no drives designed or marketed primarily for archive use.

The formula used to rank design criteria was:

((100-passes)/10)+
(absolute value of error factor/2)+
(construction 3=moderate usage, 5=high usage)+
(head contact 3=contact, 5=min contact)
/ 2.71 (to adjust the highest rank to 10)

Technology	Serpentine tracks/ Passes	Target Market	Tape Composition	Uncorrected Error Rate	Cartridge Construction Rating	Head Contact	Ranking
Sun T10000	768/24	Enterprise	Advanced MP	$1x10^{-19}$	High usage	Min contact	10.0
Sun T10000B	1152/36	Enterprise	Advanced MP	$1x10^{-19}$	High usage	Min contact	[redacted]
HP LTO4	896/56	Backup	Thin film MP	$1x10^{-17}$	Moderate usage	Contact	7.0
HP LTO5	TBD	Backup	Thin film MP	$1x10^{-17}$	Moderate usage	Contact	[redacted]
DLT-S4	1280/80	Backup	Advanced MP	$1x10^{-17}$	Moderate usage	Contact	6.1
IBM TS1120	896/56	Enterprise	Advanced MP	$1x10^{-17}$	High usage	Contact	7.7
IBM TS1130	TBD	Enterprise	Advanced MP	$1x10^{-17}$	High usage	Contact	[redacted]

Table 3-1 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 $1x10^{-17}$.

(yellow highlighted text indicates unverified information)

3.2 Transfer Rate

Transfer rate is important since it dictates 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 60 MB/sec, with 80 MB/sec desired. Much of the data archived at the USGS is raster imagery which typically lacks repeatable patterns that would compress well, therefore all transfer rates cited are native (uncompressed).

Where measured transfer rates were not available, estimated 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 3-3.

The ranking was determined by adding the actual/estimated 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.

Tape Drive Technology	Advertised/ Proposed Native Rate	Source Of Test Results	Actual/estimated Native Write Transfer Rate	% Of Adv.	Actual/estimated Native Read Transfer Rate	% Of Adv.	Ranking
Sun T10000	120 MB/sec	EROS	121.66 MB/sec	101%	121.87 MB/sec	101%	10.0
Sun T10000B	120 MB/sec	Estimated	121.66 MB/sec	101%	121.87 MB/sec	101%	[redacted]
HP LTO4	120 MB/sec	EROS	108.58 MB/sec	90.5%	108.46 MB/sec	90.4%	8.9
HP LTO5	180 MB/sec	Estimated	162.90 MB/sec	90.5%	162.72 MB/sec	90.4%	[redacted]
DLT-S4	60 MB/sec	Open-mag.com	61.00 MB/sec	102%	61.00 MB/sec	102%	5.0
IBM TS1120	104 MB/sec	Vendor	100.00 MB/sec	96.2%	100.00 MB/sec	96.2%	8.2
IBM TS1130	150 MB/sec	Estimated	144.30 MB/sec	96.2%	144.30 MB/sec	96.2%	[redacted]

Table 3-2 Transfer Rates
(yellow highlighted text indicates unverified information)

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 LTO4 at 757 GB of usable capacity per tape. The new minimum capacity requirement is 500 GB, with 800 GB or more desired. All of the reviewed technologies meet the 500 GB requirement based on their advertised capacity. Since much of the data archived is not compressible, all capacities are native (uncompressed). Where measured capacities were not available, estimated capacities are determined based on the accuracy of specified capacities of previous generations.

The capacities below presume that a gigabyte = 1,073,741,824 bytes. The ratings were determined by calculating 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 as noted in table 3-2 above.

Tape Drive Technology	Advertised/ Proposed Native Capacity	Measured/Estimated Native Capacity	% Of Advertised Capacity	Ranking
Sun T10000	500 GB	468.00 GB	93.6%	6.0
Sun T10000B	1.0 TB	936.00 GB	93.6%	
HP LTO4	800 GB	757.70 GB	94.7%	9.6
HP LTO5	1.6 TB	1.51 TB estimated	94.7% estimated	
Quantum DLT-S4	800 GB	785.00 GB	98.1%	10.0
IBM TS1120	500 GB	475.00 GB estimated	95.0% estimated	6.0
IBM TS1130	1.0 TB	950.00 GB estimated	95.0% estimated	

Table 3-3 Storage Capacities
(yellow highlighted text indicates unverified information)

3.4 Cost Analysis

Table 3-4 shows the relative drive and media costs, maintenance costs, and the cost per terabyte for media. Rankings were established by setting the cheapest (drive, maintenance, media) to 10 then rating each of the others against the lowest cost. Maintenance is based on actual annual costs after the end of warranty. 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.

Technology	Drive $/each	5 yr Maint	Drive Warranty	Media $/each	Media $/TB	Ranking Drive Cost	Ranking Maint Cost/5yr	Ranking Media Cost/TB
Sun T10000	$24,013	$6,820	12 mo	$125	$250	0.8	0.4	2.8
Sun T10000B	$24,013	$6,820	12 mo	$125	$125			
HP LTO4	$2,500	$277	36 mo	$78	$97	7.6	9.8	10.0
HP LTO5	$4,000	$277	36 mo	$130 est	$81 est			
Quantum DLT-S4	$1,910	$271	36 mo	$88	$110	10.0	10.0	8.8
IBM TS1120	$26,270	$8,736	12 mo	$161	$228	0.7	0.3	4.3
IBM TS1130	$26270	$8,736	12 mo	$161 est	$161 est			

Table 3-4 Drive, Maintenance, and Media Costs
(yellow highlighted text indicates unverified information)

3.5 Scenarios

Table 3-5 shows the total drive and media cost for three scenarios. These scenarios presume that each dataset or project stands on its own, although pooling resources for multiple datasets can mitigate cost. Note that media prices are expected to drop considerably within 6 months after product introduction. Rankings 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 system interfaces or cables.

Technology	100 TB 2 drives	200 TB 3 drives	400 TB 4 drives	100 TB Ranking
Sun T10000	$73,026	$122,039	$196,052	2.0
Sun T10000B	$60,526	$97,039	$146,052	
HP LTO4	$14,700	$26,900	$48,800	10.0
HP LTO5	$16,100	$28,200	$48,400	
Quantum DLT-S4	$14,820	$27,730	$51,640	9.9
IBM TS1120	$75,340	$124,410	$196,280	1.9
IBM TS1130	$68,640	$111,010	$169,480	

Table 3-5 Scenario Costs (drives, media)
(yellow highlighted text indicates unverified information)

3.6 Vendor analyses

Table 3-6 is intended to provide an analysis of each company and the stability of each technology. All seem to be established and stable companies, and this rating should in no way be viewed as a market analysis. When selecting an archive technology, it makes sense to look at the company and product histories though rating vendor history is challenging due to mergers and acquisitions. For SDLT and T10000, the technologies were based on predecessors (DLT and 9940); therefore, the technology age included those predecessors. The longevity rankings were determined by the following formula:

(company age + technology age) / 11.0 (to adjust the highest rank to 10)

Company	Technology	Years in business	Technology age in years	Longevity Ranking
Sun/STK	T10000	39 (1969)	8 (2000)	4.3
HP	LTO	69 (1939)	8 (2000)	7.0
Quantum	SDLT (DLT)	28 (1980)	19 (1989)	4.3
IBM	3592 (3590)	97 (1911)	13 (1995)	10.0

Table 3-6 Vendor Analyses

3.7 Drive compatibility

Table 3-7 shows the level of intergeneration drive compatibility as well as the future drives planned. The columns "% Previous Generations Read" and “% Previous Generations Written” indicate the percentage of previous generations which are read/written by the generation indicated. Drives that are the first of their generation receive a score of 50%, since it would be unfair to penalize them simply for being the first generation. 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)) / 28 (to adjust the highest rank to 10)

Technology	% Previous Generations Read	% Previous Generations Written	Future Generations Planned	Ranking
Sun T10000	50	50	5	7.1
Sun T10000B	100	100	4	
HP LTO4	66	33	2	5.0
HP LTO5	50	25	1	
Quantum DLT-S4	66	0	0	2.4
IBM TS1120	100	100	4	10.0
IBM TS1130	100	50	3	

Table 3-7 Drive Compatibility
(yellow highlighted text indicates unverified information)

3.8 Ranking summary

The ranking summary provides a quick reference to the rankings.

Drive	Design Criteria	Capacity	Media Cost	Drive Compat.	Transfer Rate	Drive Cost	5yr Maint Cost	Vendor Analyses	Scenario Cost
Sun T10000	**10.0**	**6.0**	**2.8**	**7.1**	**10.0**	**0.8**	**0.4**	**4.3**	**2.0**
HP LTO4	**7.0**	**9.6**	**10.0**	**5.0**	**8.9**	**7.6**	**9.8**	**7.0**	**10.0**
DLT-S4	**6.1**	**10.0**	**8.8**	**2.4**	**5.0**	**10.0**	**10.0**	**4.3**	**9.9**
IBM TS1120	**7.7**	**6.0**	**4.3**	**10.0**	**8.2**	**0.7**	**0.3**	**10.0**	**1.9**

Table 3-8 Ranking Summaries
(Blue indicates the highest ranking in category)

4.0 Conclusions and Recommendations for USGS Offline Archiving Requirements

4.1 Weighted Decision Matrix

The following table 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 describe 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 Bit Error Rate (BER), and amount of head contact. (See table 3-1)
- Capacity: This criterion describes the measured or estimated capacity per cartridge, which is typically less than the advertised capacity. (See table 3-3)
- Media cost/TB: This criterion is a rating of the relative cost per terabyte for media using the advertised capacity. (See table 3-4)
- 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 will give an indication of the ability to maintain drives that can read an aging archive. (See table 3-7)
- Transfer rate: This criterion describes the aggregate read and write transfer rate, which is typically less than the advertised transfer rate. (See table 3-2)
- Drive cost: This criterion is the rating of relative cost of each drive at the lowest currently available price. (See table 3-4)
- 5-year maintenance cost: This criterion rates the relative cost of maintenance over the first 5 years, taking into account warranty. (See table 3-4)
- Vendor analyses: This criterion is the rating of the viability of the vendor and technology. (See table 3-6)
- Scenario cost: This criterion is the rating of the cost of scenario #1. This includes media cost and drive cost. The measured or estimated capacity is used rather than advertised capacity. (See table 3-5)

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

Selecton Criteria	Wt	Sun T10000	HP LTO4	Quant DLT-S4	IBM TS1120	Sun T10000	HP LTO4	Quant DLT-S4	IBM TS1120
Design criteria		10.0	7.0	6.1	7.7	0.0	0.0	0.0	0.0
Capacity	20	6.0	9.6	10.0	6.0	120.0	192.0	200.0	120.0
Media cost/TB		2.8	10.0	8.8	4.3	0.0	0.0	0.0	0.0
Compatibility	15	7.1	5.0	2.4	10.0	106.5	75.0	36.0	150.0
Transfer rate	15	10.0	8.9	5.0	8.3	150.0	133.5	75.0	124.5
Drive cost		0.8	7.6	10.0	0.7	0.0	0.0	0.0	0.0
5yr maint cost		0.4	9.8	10.0	0.3	0.0	0.0	0.0	0.0
Vendor analyses	15	4.3	7.0	4.3	10.0	64.5	105.0	64.5	150.0
Scenario cost	35	2.0	10.0	9.9	1.9	70.0	350.0	346.5	66.5
Total Weighted Score						**511.0**	**855.5**	**722.0**	**611.0**

Table 4-1 Weighted Decision Matrix

4.2 Conclusions and notes

LTO4 achieved the highest total score in the study. At this time, there is no compelling reason to adopt a new standard offline archive device.

There was no opportunity to test DLT-S4 or TS1120 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 specs that were achievable in the past).

- When multiple copies of a dataset are maintained, it becomes acceptable to trade cost and performance for reliability, particularly when the working copy is on an enterprise technology such as Sun T10000, 9940, or 9840 as is the case for 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 Sun and IBM pricing, enterprise drives such as the T10000 and TS1120 are unlikely to achieve a level of market saturation that would cause significant price decreases.

- With proper handling and multiple copies, any of the technologies evaluated in this report could be deployed for archive use.

4.3 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 should be migrated to LTO4 or LTO5 in the next 2 years. Data stored on LTO3 need not be migrated in the near future unless media degradation is suspected or observed.

3. There may be a benefit in placing two or more LTO4 drives into the SL8500 robotic library to automate generation of offline copies. Once written, tapes would be ejected and stored elsewhere.

4. To reduce risk, the USGS should continue the strategy of storing datasets on multiple technologies. An example of this would be to store a working copy of a dataset on nearline T10000 and offline/offsite copies on LTO4. This strategy partially mitigates the risks of one or the other technology failing or being retired prematurely.

5. A second nearline copy on T10000 is advised, with this copy being ejected and stored locally. This would provide fast recovery without having to retrieve the offsite LTO copy.

6. 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, but it should be frequent enough to provide an opportunity to detect deteriorating media. At a minimum, a 10% annual random sampling should be implemented by all EROS projects.

7. 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 from the offline media, integrity can then be verified. Verification of each retrieved file may not be feasible due to CPU impacts.

8. It is advised to migrate all data to new media between 3 and 5 years after it was written. While 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 the National Archives and Records Administration (NARA).

9. When writing archive tapes, the tapes should be verified on a second drive. This will help identify any drive incompatibility.

10. Each time this study is revisited, it is possible that the highest scoring technology will change. This does not indicate that the USGS should change offline tape technologies frequently. There is benefit in staying with a given technology for several years, even if it is not the leading technology continuously. This study is a snapshot in time, and results would differ even a few months earlier/later due to new hardware releases. There currently is no compelling reason to abandon LTO technology.

11. The USGS should plan to update this trade study periodically so that when a logical replacement for LTO ultimately emerges, it will not be a surprise.

Abbreviations and Acronyms

AIT	Advanced Intelligent Tape
BER	Bit Error Rate
CD-ROM	Compact Disc - Read Only Memory
CERN	Conseil European pour la Recherché Nucleaire
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)
HD-DVD	High Definition Digital Versatile Disc (formerly Digital Video Disc)
HDT	High Density Tape
HP	Hewlett Packard
HW	Hardware
IBM	International Business Machines
LP DAAC	Land Processes Distributed Active Archive Center
LTO	Linear Tape Open
MB	Megabytes (1,048,576 bytes)
NARA	National Archives and Records Administration
QIC	Quarter-inch Cartridge
SAIT	Super Advanced Intelligent Tape
SD	South Dakota
SGT	Stinger Ghaffarian Technologies, Inc.
SDLT	Super Digital Linear Tape
STK	StorageTek (now a Sun business unit)
TB	Terabytes (1,024 GB or 1,099,511,627,776 bytes)
TBD	To Be Decided/Determined
USGS	United States Geological Survey

Citations

This section was an afterthought and therefore the body of the study does not refer to specific citations listed here. A majority of my references are listed here, but not all of the citations listed here were specifically used in the study. Where I used magazine articles, I've listed a link to the online copy.

Vendor sites:

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

http://www.sun.com/storagetek/tape_storage/tape_drives/ (Sun/STK)

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

CERN reports:

http://cscct.home.cern.ch/cscct/LTO3.ppt

http://cscct.home.cern.ch/cscct/T10000.ppt

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.fcw.com/article94344-05-15-06-Print&newsletter%3Dyes

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

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

http://www.infostor.com/articles/article_display.cfm?Section=ARCHI&C=Labsr&ARTICLE_ID=255609&KEYWORDS=super%20tape%20paradox&p=23

http://au.sun.com/edge/pdf/t10000drive_technicalbrief.pdf

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

http://www.inphase-technologies.com/products/default.asp?tnn=3

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

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-03.ibm.com/systems/storage/tape/ts1120/index.html

http://www.sun.com/storagetek/tape_storage/tape_drives/t10000/488126_T10Kdrive_WP_LF.pdf http://www-03.ibm.com/systems/hu/pdf/file_hu_prztte01.pdf

http://ch.sun.com/sunnews/events/2006/apr/solutionday/pdf/5_Welche_Loesungen_bringt_Sun_im_Bereich_Mainframe_Tape.pdf

http://h18006.www1.hp.com/products/storageworks/ultrium1840/index.html

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

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