

ESG Lab Review

LTO for the Next-generation Data Center


Date: November 2015 **Author:** Vinny Choinski, Senior Lab Analyst; and Brian Garrett, VP ESG Lab

Abstract: This ESG Lab Review is designed to address any possible misconception that LTO tape is a slow or unreliable storage platform, while also exploring its new features and use cases including the Linear Tape File System (LTFS), tape-as-NAS, deep and active archiving, and consolidated tape-out data protection solutions.

The Challenges

Two responses that are continually cited among the top five IT priorities identified by respondents in ESG's annual IT spending intentions survey are improving data backup and recovery and managing data growth (see Figure 1).¹ It's not a coincidence that these responses often appear together in ESG survey results since data growth has a direct impact on an organization's ability to back up and restore data in a timely manner.

Figure 1. Top Five IT Priorities



Source: Enterprise Strategy Group, 2015.

Historically, data growth in a production environment meant adding more disk to the primary storage environment, which in turn meant adding more tape or secondary disk storage to the backup environment. Now, this dynamic is changing, as new enhancements in LTO Ultrium technology are blurring primary and secondary storage lines and giving customers more choices in how they store and protect their business-critical data.

¹ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by The LTO Program, including IBM, HP, Quantum, and Fujifilm.

The Solution: Linear Tape-Open (LTO) Deployments

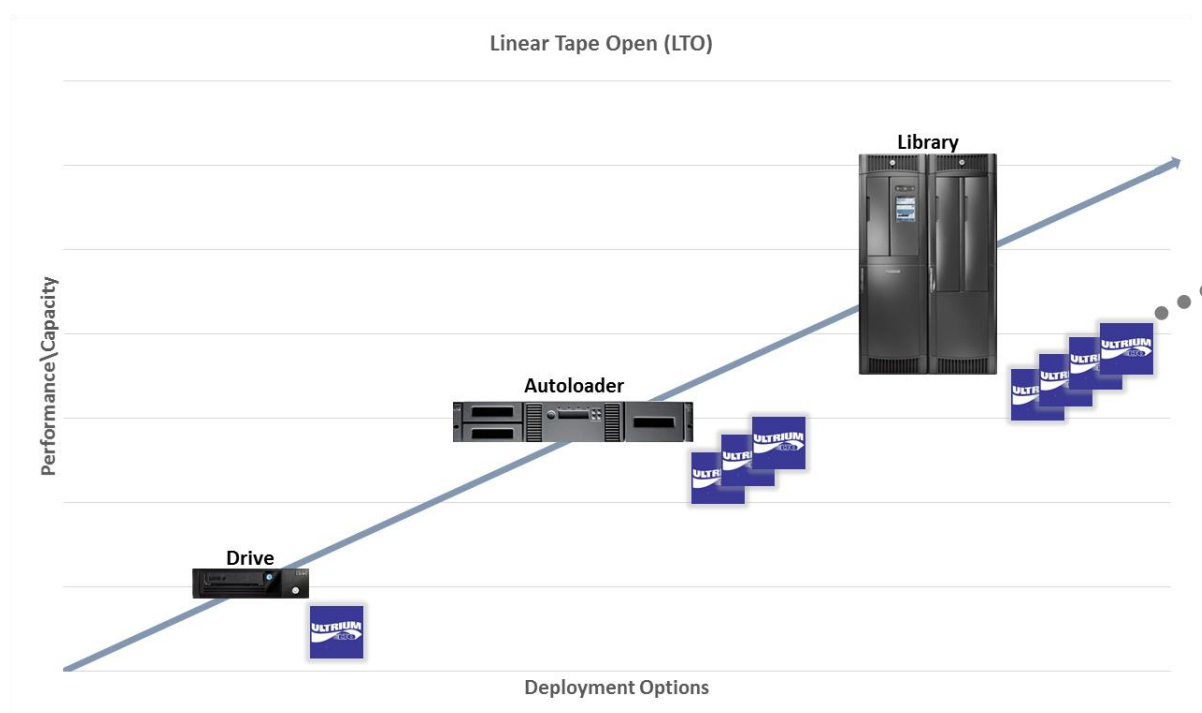
Originally released in 2000, the Linear Tape-Open (LTO) Ultrium technology was developed as an alternative to proprietary tape formats. Development, licensing, and certification of media and drives are governed by the LTO Consortium, which maintains rigid interoperability standards so that drives and cartridges from different vendors are expected to be interchangeable. Now in its seventh generation of availability, LTO Ultrium high-capacity magnetic tape storage technology is capable of storing 15 TB of data on a single cartridge with up to 750MB/s transfer rate.² Key metrics include:

- A development roadmap for the technology with defined capacity capabilities all the way out to LTO-10.
- An algorithm that delivers 2.5:1 lossless compression for LTO-6 and LTO-7 and 2:1 compression for older generations.
- Compatibility rules that define Ultrium technology interoperability:
 - A drive can read data from media in its own generation and two prior generations.
 - A drive can write data to media in its own generation and the immediate prior generation.

Initially, the technology was predominantly used as a target in data protection solutions, but the introduction of the Linear Tape File System (LTFS) in generation five expanded the scope of LTO in the modern data center significantly. Now, organizations are leveraging it to address the challenges of managing ever-growing tertiary and even primary storage. As shown in Figure 2, LTO technology can be deployed in diverse configurations from small standalone drive environments, to mid-market autoloader and library solutions, all the way to multi-frame enterprise library infrastructures. Next-generation data center enablement features include:

- Write Once Read Many (WORM) technology introduced in LTO-3. This is typically used for legal record keeping.
- Application-managed, natively-enabled tape drive level AES-GCM encryption introduced in LTO-4.
- Partitioning introduced in LTO-5, which enabled the tape to be divided into separately writable areas.
- Linear Tape File System (LTFS), a self-describing tape file system made possible by the partitioning feature.

Figure 2. Solution Deployment Options



² Capacity and data transfer rates assume 2.5:1 compression and optimal transfer conditions.

ESG Lab Review

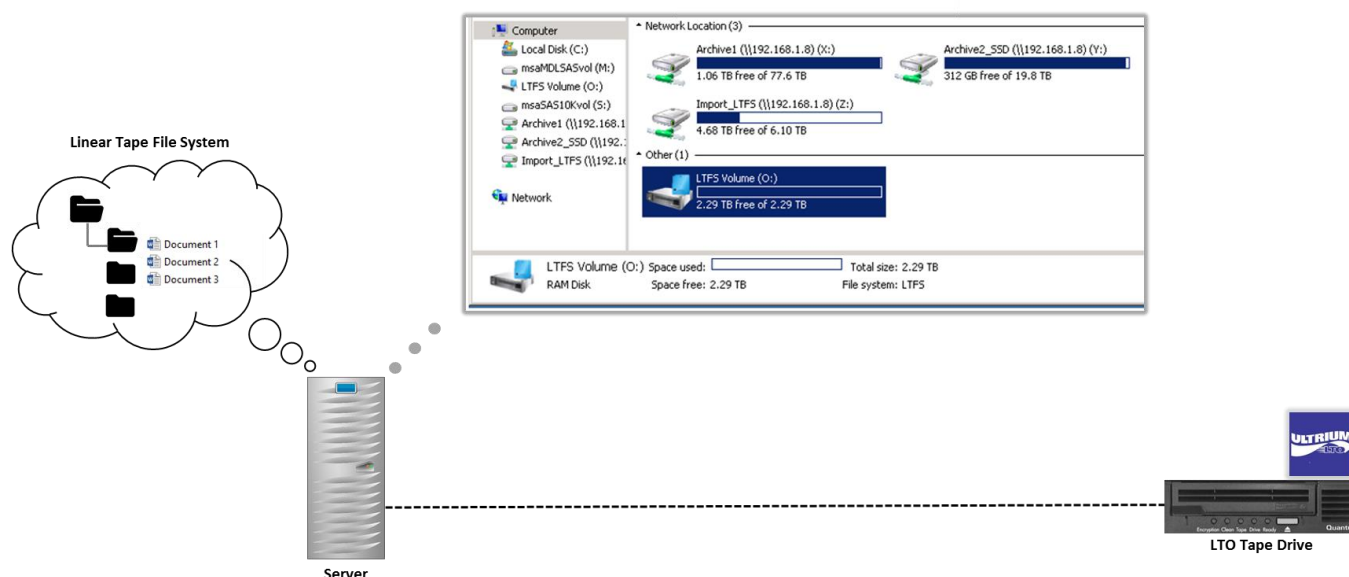
ESG Lab performed hands-on evaluation and auditing of LTO Ultrium technology at a Hewlett Packard Enterprise facility in Fort Collins, CO. Testing was designed to demonstrate reliability, improved performance and capacity, and new features that move the technology beyond only data protection and enable its use in multiple storage configurations.

Modern Use Case

ESG Lab started its validation of LTO technology by reviewing how newer generation LTO features are being leveraged in today's data center. The introduction of partitioning, in LTO generation five, enabled many of these modern use cases. Partitioning allows formatting the tape into one section that holds content or data and another section that holds an index of that content, thereby creating self-describing media. Partitioning made the Linear Tape File System (LTFS) possible. LTFS works in conjunction with LTO drive and media technology to make the device look like a typical storage device to the underlying operating system. Once configured, files can be simply copied to and from the device just like with any disk drive.

As shown in Figure 3, ESG Lab deployed a simple LTFS configuration on a Windows server with a SAS connected LTO-6 standalone tape drive. We simply downloaded and installed the free LTO-6 driver and LTFS utilities and then configured and mounted the tape drive as a 2.29TB volume on the test environment Windows server. Best practices recommend moving files that require modification to a random access device while using this configuration as a local or portable archive type solution and for large sequential read operations.

Figure 3. Linear Tape File System (LTFS) Standalone

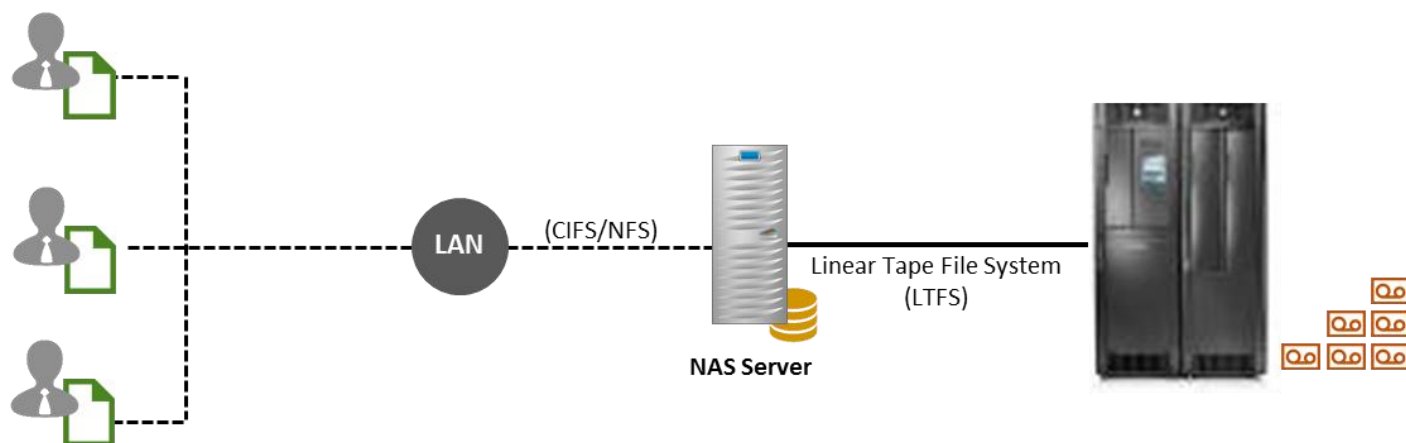


Free downloadable LTFS solutions include:

- Hewlett Packard Enterprise StoreOpen Standalone and Automation.
- IBM Linear Tape File System Single Drive Edition.
- Quantum Linear Tape File System.

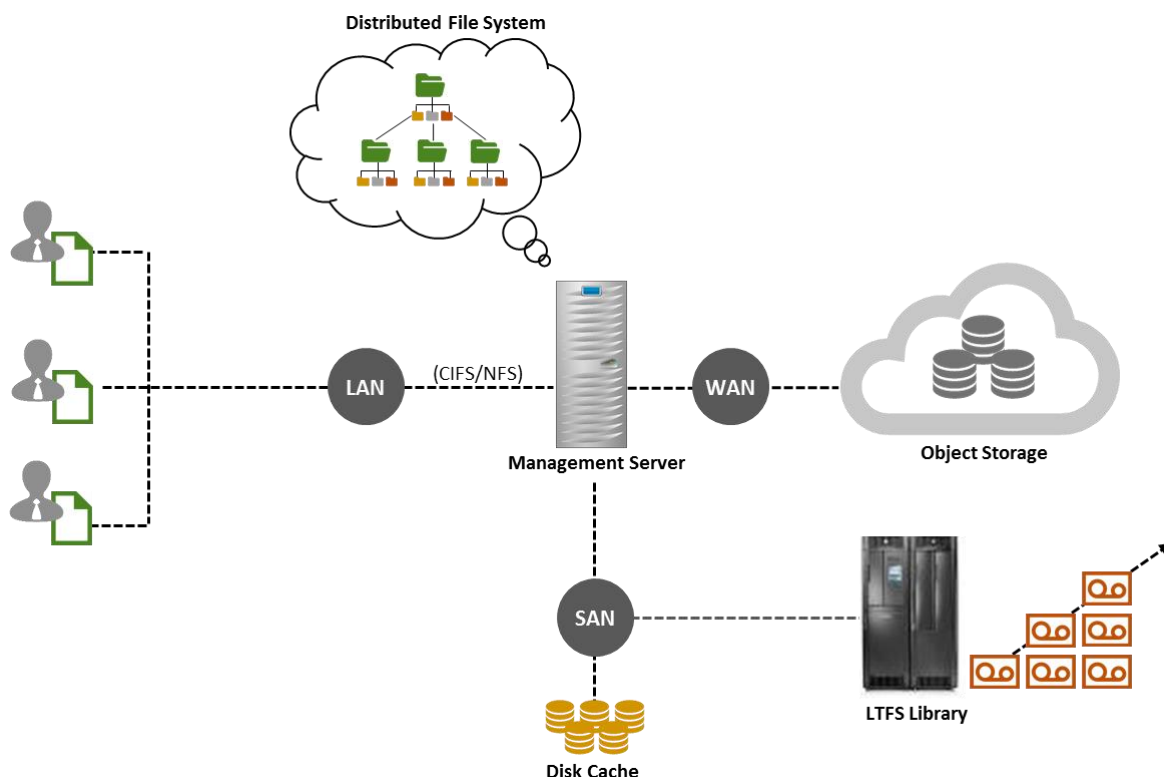
Next, ESG Lab explored expanding the local capabilities of LTFS to a NAS solution. Here, LTFS-formatted media is stored in an automated LTO tape library and is used as high-capacity, cost-efficient back-end storage to a front-end NAS server. The NAS server manages the presentation of data to the end-users over the corporate LAN via the CIFS/NFS protocols. It also manages the movement of active data between LTFS storage and disk cache. This process enables active data (such as those files currently being modified) to be serviced from a small amount of disk cache on the NAS server, and less active data to be stored on large-capacity LTO media in the automated library as an active archive.

Figure 4. Tape-as-NAS




Next, as shown in Figure 5, we explored a tiered storage architecture with tape, disk, and cloud storage. Here, data is still presented to end-users via the tape-as-NAS solution. Then, a policy engine moves data based on business requirements to the most efficient storage media, disk cache, tape, and cloud.

Figure 5. Distributed Storage Environment



Finally, ESG Lab reviewed best practices for tape in today's data protection solutions. Even with the adoption of disk and cloud, tape remains a major component in many data protection environments. It provides customers with a fast, highly reliable, and portable storage media. A well-designed and balanced data protection architecture often leverages tape library solutions for active archive and consolidated output. Large sequential data streams and groups of disk backups with high throughput requirements can be staged to the library's portable media for efficient offsite data transfer. Less bandwidth-intensive workloads are sent directly to the cloud.

Figure 6. Data Protection



Why This Matters

Tape technology has a long history in data protection, and even with the adoption of disk and cloud targets, it still plays a major role in many backup and archive environments. According to ESG research, 49% of backup environments leverage tape today with 23% relying solely on tape.³ Clearly tape is not a new concept for IT professionals, but some of its new features and best practices might be.

ESG Lab confirmed that the addition of partitioning and LTFS adds compelling new uses cases for LTO technology. Now, this removable media can be easily used with its native utilities for file-level operations, and these file-level capabilities can be expanded to tape-as-NAS and tiered storage solutions. Also, when deployed with modern data protection best practices, tape—combined with disk and cloud components—can extend efficiency while providing customers with huge data storage flexibility.

³ Source: ESG Research Report, [2015 Trends in Data Protection Modernization](#), September 2015.

Reliability


When it comes to reliability in tape-based data storage solutions, the number one priority should be the ability to access data when needed. A tape drive may fail, but if the data can still be read back with another drive, then retrieval is successful. Other important considerations include media durability, media shelf life, drive reliability, and probably most importantly, end-user best practices.

ESG Lab began its review of LTO solution reliability by interviewing members of the quality assurance team from an LTO Program participant, along with exploring LTO media characteristics. LTO media was designed on the premise that it would be loaded into a device, written to, removed, moved, stored, recalled, and reloaded into a device when data needed to be retrieved. With that in mind, each consortium member is constantly working to keep LTO media highly reliable, portable, and rugged enough to be moved without impacting data integrity.

LTO Ultrium media reliability starts with format design. First, data is spread across multiple channels (8, 16, or 32, depending on generation) to protect against single-head element or media defects. Then error detection/correction code (ECC) is applied at two levels. Level one ECC rewrites the data to another location further down the original track. Level two ECC rewrites the data across multiple tracks and allows for data recovery in the event of a totally bad track. The process uses read-while-write verification where the read head checks the validity of each dataset written. If a data sub-set logs an error, it is automatically rewritten. These format design features produce a one in 10^{19} bit-error-rate (BER) for LTO-7 media. In layman's terms, this means that it would take 130 tape drives writing data continually for one year to encounter an error that could not be fixed by ECC. As shown in Figure 7, you are more likely (1 in 10^{16}) to hit an uncorrectable error in your enterprise disk environment. In fact, you would have a much higher probability of:

- Getting hit by lightning; the odds are one in a million.
- Getting killed by a shark; the odds are one in 11.5 million.
- Winning a multi-million dollar lottery; the odds are 1 in 259 million.

Figure 7. Media Error Rate




Next, ESG Lab reviewed the product testing methodology used to ensure the reliability of the tape drives that are used to read and write data on LTO Ultrium media. For this testing, three high-level concepts were used:

- Mean time between failures (MTBF) is used to show that the product will achieve expected reliability.
- Design verification covers a full suite of ambient, environmental, and regulatory compliance tests.
- Field warranty and repair data is collected and monitored to validate annual failure rates.

The result of an environmental LTO-6 Constant Transfer Rate test is displayed in Figure 8. This is just one of the many rigorous tests ESG Lab audited during the review to validate the commitment of LTO Consortium members to delivering highly reliable solutions. It shows the transfer rate achieved when writing full volumes to multiple green cartridges, a process often used in archiving solutions.⁴ Testing results are expressed in terms of the difference from the theoretical maximum native value of 160 MB/sec for LTO-6 tape drives. It should be noted that, for the duration of the test, the drive maintained extremely consistent throughput (e.g., after writing 2,000 green tape cartridges).

ESG auditing of annual intervention rates (e.g., service calls) indicates that reliability issues are far more likely to stem from system hardware or software issues, operator error, or improper media handling and storage than from an LTO Ultrium drive error. This highlights the importance of implementing end-user best practices for any LTO Ultrium solution.

Figure 8. Transfer Rate Reliability Testing



Why This Matters

ESG Lab has validated through extensive auditing of reliability test results and review of the rigorous testing procedures that buyers of LTO Ultrium technology should be extremely confident about reliability at the device or media level. In fact, users are far more likely to encounter data loss with enterprise disk than with LTO tape. ESG Lab confirmed that the most important factor for deploying a highly reliable solution is a clear focus on well-documented best practices when implementing the hardware infrastructure and software applications that leverage this technology.

⁴ Green tape refers to new LTO media, the most abrasive condition to the LTO drive tape head.

Performance and Scale

To validate performance and scale, ESG Lab reviewed the generation over generation capacity and transfer rate improvements achieved at the individual component level (e.g., tape drive and cartridge) as well as the ability to scale a solution via multiple drives and cartridges in automated tape library configurations.

Figure 9 shows the increase in native capacity and transfer rates of the technology from the original release of LTO-1 in 2000 through the newly released LTO-7 version. Native capacity increased from 100 GB in generation one to 6,000 GB in generation seven. Native transfer rates increased from 20 MB/s to 300 MB/s respectively.

Figure 9. Performance and Capacity




Table 1. Compressed Capacity and Transfer Rates for LTO-3 through LTO-10

Generation	Capacity	Transfer Rate
LTO-3	800 GB (2:1 Compression)	Up to 160 MB/s (with 2:1 Compression)
LTO-4	1.6 TB (2:1 Compression)	Up to 240 MB/s (with 2:1 Compression)
LTO-5	3 TB (2:1 Compression)	Up to 280 MB/s (with 2:1 Compression)
LTO-6	6.25 TB (2.5:1 Compression)	Up to 400 MB/s (with 2.5:1 Compression)
LTO-7	Up to 15 TB (2.5:1 Compression)	Up to 750 MB/s (with 2.5:1 Compression)
LTO-8	Up to 32 TB (with 2.5:1 Compression)	Up to 1,180 MB/s (with 2.5:1 Compression)
LTO-9	Up to 62.5 TB (with 2.5:1 Compression)	Up to 1,170 MB/s (with 2.5:1 Compression)
LTO-10	Up to 120 TB (with 2.5:1 Compression)	Up to 2,750 MB/s (with 2.5:1 Compression)

What the Numbers Mean


- Compressed transfer rates of LTO-3 (160 MB/s) are faster than a typical 7,200 RPM disk drive (130 MB/s).
- Performance and capacity are improving significantly generation over generation.
- The italicized data in Table 1 represents LTO roadmap metrics for generation seven through ten.⁵

⁵ <http://www.lto.org/technology/what-is-lto-technology/>

Finally, ESG Lab explored the scalability aspects of an LTO solution and how customers are leveraging the technology to address the challenges of managing very large storage environments. Industries such as media and entertainment, Internet, telecommunications, bio-tech, oil and gas, aero-space, finance, and public and private research organizations often turn to tape-based solutions to solve these challenges.

As shown in Figure 10, one of the largest super computer environments used for multiple research programs and applications—including hurricane/tornado modeling, the Big Bang, and creating a “computational microscope” for bio-molecular structure and modeling studies—leverages LTO storage. The solution has a storage capacity of 380 petabytes (PB)—the equivalent of 5,054 years of HDTV video or a stack of books reaching over nine times the distance from the Earth to the moon—and can reach read/write rates that exceed 1 terabyte/sec, settling at about 2.2 PB/hour.

Figure 10. Solution Scale



Why This Matters

The increased use of virtualization, ever growing amounts of data that need to be managed, and initiatives for improving backup/recovery and archival environments all put pressure on IT professionals as they design and implement today's storage solutions. What's needed is a solution with the flexibility and agility to deliver high performance and scalable capacity as storage demands grow and environments change.

ESG Lab has validated that LTO Ultrium technology can help organizations meet these data storage challenges. With single drive compressed data transfer rates of up to 750 MB/s, compressed storage capacity of 15 TB per cartridge for LTO-7, and the ability to get tens, or even hundreds, of tape drives working in parallel, ESG Lab confirmed that an LTO solution can be scaled to meet just about any storage needs.

The Bigger Truth

Contrary to any misconception that tape is unreliable, slow, and simply overlooked as a modern storage technology, ESG is pleased to report that LTO tape is thriving and has a bright future in organizations of all sizes, including some of the largest public cloud providers on the planet. An ESG Lab audit of LTO consortium reliability testing has confirmed that LTO-7 media and drives are orders of magnitude more reliable than enterprise-class disk drives. As a matter of fact, it's more than a billion times more likely that you will get hit by lightning in your lifetime than that you will encounter an unrecoverable LTO-7 media error.

Also, ESG's annual IT spending intentions research survey indicates that the top three considerations reported by respondents for justifying IT spending in 2015 were security/risk management, cost reduction, and data analytics initiatives that allow business teams to extract incremental value out of their existing information repositories.⁶ With these considerations in mind, it's clear to ESG that modern tape with industry-standard LTO can play a key role as organizations transform their IT architectures based on economics and agility. LTO tape continues to play a key role in enterprises as it reduces risk for the organization with field-proven backup/recovery and archive technology that's more cost-effective and reliable than disk. Recent advances in the evolution of tape technology (e.g., the industry-standard LTFS and partitioning support that's built into LTO) are enabling a new wave of tape modernization and online active archive solutions that can drastically reduce the cost and complexity of big data initiatives, including creating higher margins for both its on-premises and cloud storage providers.

ESG Lab has confirmed that LTO technology can be used to meet the performance needs of the most demanding streaming workloads (e.g., real-time geospatial data archiving and high-definition video production) with native, uncompressed, single tape drive LTO-7 performance of up to 300 MB/sec and the ability to get tens, or even hundreds, of tape drives working in parallel in one or more enterprise-class tape libraries.

While ESG confirmed the reliability of LTO media and drives, it's not just the media and the drives that ensure data reliability. Human factors and software errors are the primary risk factors that can lead to data loss with tape or any other technology. Following time-tested industry best practices for system diagnostics, data protection and retention, and tape media handling and storage reduces that risk.

While ESG believes that the reliability and the tape modernization analysis presented in this report are vital considerations when making a long-term data retention technology decision, the bottom line for most IT professionals is simple economics. After all, if economics weren't an issue, most organizations would store all of their application data forever on disk or in the cloud! Nothing is more cost-effective, reliable, and energy-efficient for long-term data retention than a tape in a library slot or on a shelf, and it continues to play a key role for organizations across the globe. Isn't it about time we embrace the renaissance of tape and leverage the innovations that make it ever more useful?

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

⁶ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.