

ESG Lab Spotlight

Efficient, Reliable Archive from Hewlett Packard Enterprise

LTO for the Next-generation Data Center

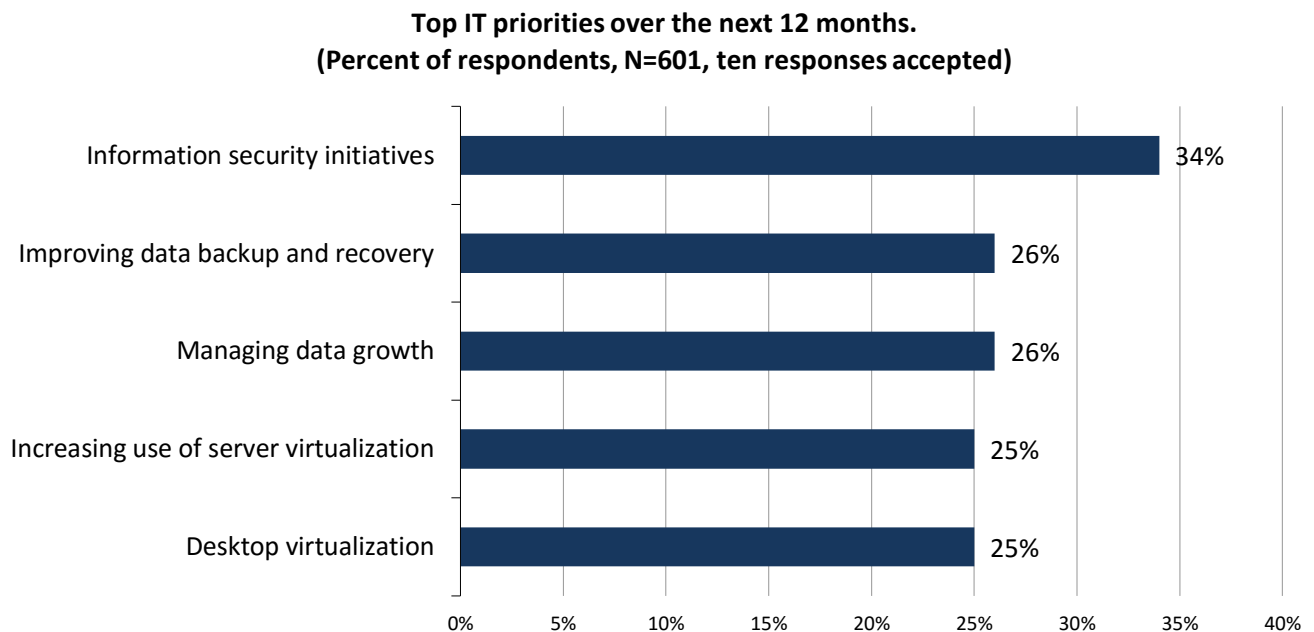
Date: June 2016 **Author:** Jason Buffington, Principal Analyst; and Vinny Choinski, Senior Lab Analyst

Abstract: This ESG Lab Review excerpt and Spotlight is designed to address any possible misconception that LTO tape is a slow or unreliable storage platform, while also exploring new features and use cases of HPE's StoreEver Archive Manager.

The Challenges

Two responses that are continually cited among the top five IT priorities identified by respondents to 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 Spending Priorities in 2015



Source: Enterprise Strategy Group, 2016.

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, protect, and manage 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, HPE, Quantum, and Fujifilm.


Reliability and Performance

Without a doubt, the number one priority should be the ability to access data when needed. When it comes to reliability in tape-based data storage solutions, tape drive mechanics may fail, but if the cartridge can still be read back with another drive, then data retrieval is successful. Other important considerations include media durability, media shelf life, drive reliability, and probably most important, end-user best practices.

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. Next, 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 a single error that could not be fixed by ECC. As shown in Figure 7, you are more likely (1 in 10^{16}) to find 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 2. Probability of LTO Media Error Compared with Other Rare Events



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). 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 3. Performance and Capacity of LTO Solutions



What the Numbers Mean

- Compressed transfer rates of LTO-6 (160 MB/s) are faster than a typical 7,200 RPM disk drive (130 MB/s).
- Performance and capacity continue to improve significantly generation over generation.²

Why This Matters

The increased use of virtualization and other sprawling production platforms, 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 reliable solution with the flexibility and agility to deliver high performance and scalable capacity as storage demands grow and environments change.

Decades ago, those burgeoning IT issues drove many to abandon yesteryear’s tapes, under the presumption that legacy tape systems were slow and unreliable. That story is perpetuated even today by those who want to sell disk-systems exclusively. The scale demands of tomorrow are forcing a re-assessment of these presumptions, with many organizations re-discovering modern tape for what it is—*modern*.

ESG Lab has validated, through extensive auditing of reliability test results and a review of rigorous testing procedures, that buyers of LTO Ultrium technology should be extremely confident about reliability at the device and media levels. In fact, users are much more likely to encounter data loss with enterprise disk than with LTO tape. ESG Lab has also 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.

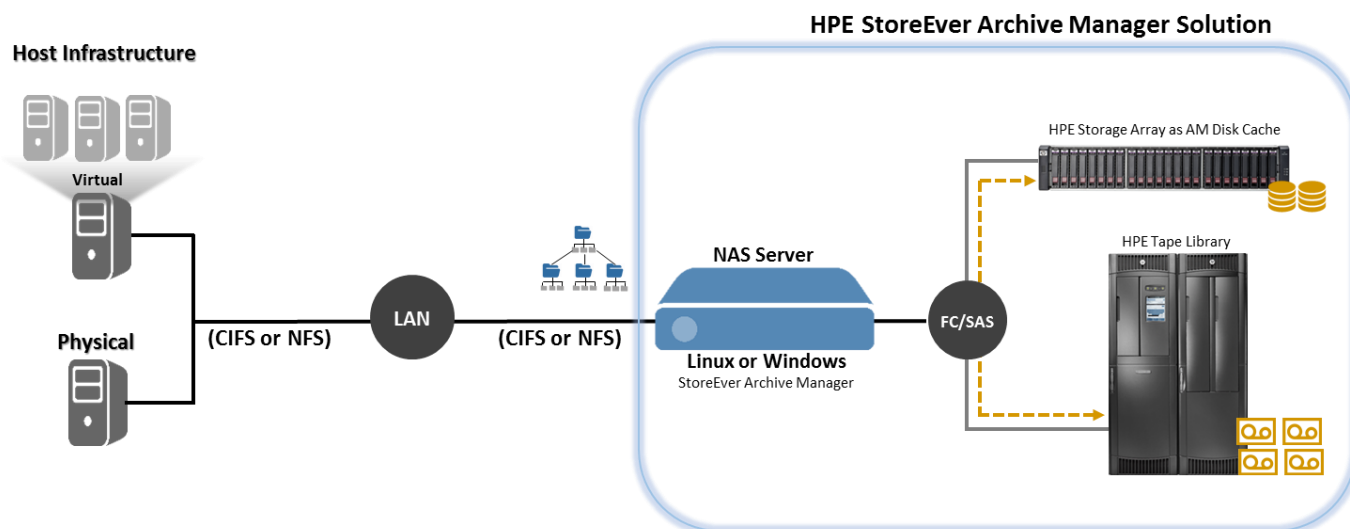
² <http://www.lto.org/technology/what-is-lto-technology/>

One Solution to Consider: StoreEver Archive Manager from Hewlett Packard Enterprise

While the previously summarized ESG Lab Report looked at the merits of LTO tape without regard to a specific vendor’s solution offerings, this section takes a closer look at a particular storage hardware and software solution that leverages LTO tape technology. Specifically, StoreEver Archive Manager from [Hewlett Packard Enterprise \(HPE\)](#).

As shown in Figure 4, the solution leverages HPE StoreEver LTO tape libraries, drives, and media for cost-efficient, high-capacity data storage; HPE disk storage arrays for caching; and HPE StoreEver Archive Manager software for data access, placement, and movement. It supports Linear Tape File System (LTFS) which provides two major advantages over proprietary, often software-driven, tape formats: 1) data portability, and 2) no vendor lock-in. This means that data stored on LTO tape media can be moved between locations where the data can be accessed by any system with a supported device and LTFS software driver installed. Archive Manager does also support a proprietary tape format that can be used in certain situations when LTFS cannot be used. As shown in Figure 4, clients access the StoreEver Archive Manager data the same way they would in a typical NAS environment. Windows systems simply map the data via CIFS/SMB while Linux systems mount the data via NFS. The solution can be easily scaled by simply adding more tape cartridges.

Figure 4. HPE StoreEver Archive Manager Solution Overview




To the end user, data is presented through the operating system interface as standard directories, sub-directories, and files—even though the data is actually stored on a series of tapes. The back-end data movement between tape and disk buffer is abstracted from the user via the StoreEver Archive Manager software. For write operations (e.g., file creation), the data is first written to a disk-based buffer and then copied to tape. If a file is modified, Archive Manager creates a new version of this file and stores it like a new file. Users can “roll back” to a specific file version at any time, even if the original file is deleted. For read operations, the size of the disk buffer will determine if the data is read from disk or tape. Older content in most cases will be read from tape, while more active data will be read from the disk buffer.

The HPE StoreEver Archive Manager solution has multiple levels of integrity checking, including checksums and CRC validation at the file and media levels, and automatic replication to protect against data loss. As shown in Figure 5, data integrity visibility is provided through the HPE Command View for Tape Libraries software. HPE StoreEver Command View for Tape Libraries is a software application that enables management, monitoring, and configuration of all HPE tape libraries through a single console. It provides remote management, diagnostics, and configuration of HPE StoreEver tape libraries across multiple locations. Figure 5 shows the results of data verification within Command View.

To the end-user, data is presented through the operating system interface as standard directories, sub-directories, and files—even though the data is actually stored on a series of tapes.

The data verification process can be used to ensure data integrity for less active content in the HPE StoreEver Archive Manager environment by creating a verification partition in the tape library. This view provides descriptive information about the media, its reliability, and the status of the verification test, as shown in the red call-out box on the right side of the figure.

Figure 5. HPE StoreEver Command View for Tape Libraries



Why This Matters

Organizations often choose disk over tape because they perceive it to be more reliable. The argument is that data is rarely written to single spindles, but instead implements supplemental error correction mechanisms such as RAID and software-driven, controller-based error handling. These approaches are essential for protection storage, given that these are the copies of last resort. So it makes sense that HPE implemented supplemental error handling in this tape-enabled solution.

By combining a NAS front-end controller that is disk-based within the solution, HPE addresses one of the last potential drawbacks of a large tape-based storage farm: accessibility for random requests. While individual LTO cartridges are more than fast enough for disparate read/write access, the loading/unloading of cartridges within a library might still be considered inaccessible to client devices that are leveraging NAS protocols and presume NAS (disk)-speed. Utilizing a disk-based cache mitigates that challenge, thus providing the responsiveness within the CIFS/NFS protocols, while still taking advantage of the lower cost and higher capacity of tape.

ESG Lab verified NAS functionality with no perceptible performance differences for most scenarios over what a user accessing a traditional disk-based NAS might experience. HPE’s understanding of random production/user access and how most backup and archival software handles streaming to a wide variety of devices with differing throughput behaviors enables the solution to deliver similar results to disk-based accessibility with lower overall storage costs.

The Bigger Truth

Contrary to the 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 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 survey indicates that the top two considerations reported by respondents for justifying IT spending in 2015 were *security/risk management* and *cost reduction*.³ When asked how those organizations would control costs, *investing in new technology with higher ROI* was among the top responses. 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 serve an important function in enterprises as it reduces risk for the organizations, leveraging 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 on-premises implementations and cloud storage providers.

Primary NAS solutions are evolving such that a flash-based caching tier is almost the default configuration, with the colder data being stored on less expensive, less performant, but greater capacity disk spindles. But using that rationale, where immediate performance isn't the solution driver (due to a caching tier), tape is even less expensive for high capacity, while still more than "fast enough" to retrieve that data which is not in the cache. One could make the argument that HPE StoreEver Archive Manager foreshadows a more scalable future for NAS platforms, where flash is in the "front," and tape is in the "back"—with disk spindles actually becoming irrelevant. Considering HPE's breadth of server and storage offerings, including its leadership stance within LTO, it may be particularly well suited to help fulfill that vision.

If you would like to learn more, read the full lab report at <http://www.esg-global.com/lab-reports/> and see more about the HPE StoreEver Archive Manager solution at <https://www.hpe.com/storage/storeever>.

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.

³ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.