

Tape Has Become a Strategic Imperative in the Mass Storage Landscape

Unlocking the Value of Data Pushes Tape to Center Stage



December 2020.

The [Tape Storage Council](#), which includes representatives of BDT, Cozaint, Detron, Frontier BV, FUJIFILM, GazillaByte, Hewlett Packard Enterprise, IBM, Imagine Products, Insurgo Media, Iron Mountain, Park Place Technologies, Oracle, Overland-Tandberg, PoINT Software & Systems, Qualstar, Quantum, REB Storage Systems, Spectra Logic, StrongBox Data Solutions, StorageDNA, SullivanStrickler, Turtle, XenData and XpresspaX has issued this report to highlight the current trends, usages and technology innovations occurring within the tape storage industry.

Introduction

The year 2020 will certainly be remembered for the COVID-19 pandemic, the defining global health crisis of our time. The economic uncertainty caused by the pandemic has touched everyone and the long-term impact remains unclear. What is clear as we enter into the post COVID-19 digital economy will be the need for cost containment under increasing IT infrastructure and storage budget constraints. There will be a renewed focus on cost-effectively managing pre-COVID exponential data growth plus the influx of new data demand as organizations have shifted to all digital, remote and virtual work environments including government, education, entertainment, health care and nearly everything else. Cybercrime and ransomware attacks will continue to threaten the economic viability of large organizations as well as individuals. The continued increasing demand for IT services will result in an explosion in energy use and carbon emissions raising concerns about how to keep the industry's environmental impact and its impact on climate change in check.

The continuing growth of digital data and the need to preserve more diverse data types are also changing the storage landscape. Data is now being generated faster than it can be analyzed, significantly extending data retention timeframes. The archival usage model of storing and protecting vast amounts of data for indefinite periods of time is quickly evolving. Fortunately, modern data tape continues to be fueled by significant technological and architectural developments and this trend shows no signs of letting up. Steady advances reinforce tape's ability to continue to deliver the lowest cost, highest capacity, fastest data transfer rates, and most reliable digital storage medium available, with the reliability levels three orders of magnitude better than the best HDDs making it the most cost-effective, most energy-efficient long-term storage solution available.



2020 Highlights Increased LTO Capacity and Throughput and a New LTO Roadmap

LTO-9 Arrives

On Sept. 9, 2020 the [LTO Program](#) technology provider companies, Hewlett Packard Enterprise, IBM Corp. and Quantum Corp., announced the specs of LTO Ultrium format Gen 9 (LTO-9), which were made available for licensing. LTO-9 capacity will have 18 TB native and 45 TB compressed per cartridge, a 50% increase in capacity over LTO-8. LTO-9 delivers a 2.5:1 compression ratio, available since LTO-6, with backward read and write compatibility with the previous generation of LTO-8 cartridges. A SAS interface support at 12 Gb/sec has been added for LTO-9 and the BER (Bit Error Rate) has been improved from 1×10^{19} to 1×10^{20} .

LTO media has doubled in capacity approximately every two to three years since LTO-1 with 100 GB having arrived 20 years ago. For LTO-9, the LTO Program elected to balance cost, performance and capacity by offering an 18 TB tape cartridge to address the current market for storage demand. A new LTO roadmap has been established with the goal to double capacity in each gen moving forward. One 18 TB LTO-9 cartridge can hold 61.2 years of video recording running 24 hours per day, 4.78 billion human genomes worth of sequence information or 2.88 years of data transmissions from the Hubble Space telescope and has 180x more native capacity than LTO-1.


Enterprise Tape Adds New Interface

The TS1160 tape drive from IBM is the latest generation of enterprise class tape with a 20 TB native capacity (up to 60 TB with 3x compression) while delivering a data rate of 400 MB/sec. The TS1160 recently added SAS interface support to its robust connectivity family of offerings including 8 Gb/s Fibre Channel, 16Gb/s Fibre Channel and 10GbE or 25GbE RoCE Ethernet options.

Media Development Ongoing While Record Capacity Shipments Achieved

Tape media manufacturers continue to leverage Barium Ferrite magnetic particles while new magnetic particles are under development for future generations of tape using Strontium Ferrite (SrFe) that has the potential to store beyond 400 TB native per cartridge, or more than 22 times greater than LTO-9 capacity.


The LTO Program technology provider companies released their annual tape media shipment report, detailing annual shipments. Capacity shipments rose to record amounts in 2019. “More than 225 million LTO cartridges and more than 4.4 million drives have shipped since its introduction,” according to the LTO Program.



*Assumes a 2.5:1 compression with larger compression history buffer available beginning with LTO-6. The report showed a record 114,079 PB of total LTO tape capacity (compressed) shipped in 2019. Aggregate capacities in 2018 and 2019 do not include the enhanced capacity of LTO-7 Type M shipments. Enterprise tape capacity shipments are not included in this chart.

Tape Roadmaps - LTO Roadmap Extended to Generation 12

The latest [LTO technology roadmap](#) defines specifications through twelve (12) generations of LTO tape, extending the total capacity of data held on one LTO-12 tape cartridge up to 360 TB with 2.5x compression – an increase of 8 times the compressed capacity of LTO-9 cartridges. The LTO roadmap projects that native capacities of LTO drives will approximately double with every subsequent generation through generation 12.



NOTE: Compressed capacity for generation 5 assumes 2:1 compression. Compressed capacities for generations 6-12 assume 2.5:1 compression (achieved with larger compression history buffer).

SOURCE: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only. Linear Tape-Open, LTO, the LTO logo, Ultrium, and the Ultrium logo are registered trademarks of Hewlett Packard Enterprise, International Business Machines Corporation and Quantum Corporation in the US and other countries.

The INSIC 2015-2025 Magnetic Tape Storage Roadmap Projects Strong Future Growth

The [INSIC](#) roadmap indicates the current areal density scaling rate of HDDs to be about 16% CAGR and tape to be at 33% CAGR indicating the current cost advantage of tape systems over HDDs will grow wider in the future. The combination of BaFe (Barium Ferrite) and SrFe (Strontium Ferrite) particulate media with TMR (Tunneling Magnetoresistive) recording heads has emerged as the technology combination that will most likely carry tape through the next several generations. The areal density of tape is much lower than HDD as tape achieves its capacity advantage over hard disk by having a much larger total recording surface area in a cartridge, with about 1000X the area of a 3 1/2 inch disk platter. Therefore, tape doesn't need as high of an areal density to continue its significant cost per TB advantage over HDDs.

Tape's Growing Role in Hyperscale, Hyperscale Lite and Cloud Environments

Hyperscale, Hyperscale Lite (soon to be hyperscale) and cloud environments are the most trending and highly adopted data storage platforms. They offer infrastructure services on an extremely large scale, constantly pressuring the limits of existing IT infrastructures. To deliver their services, these companies had to find new ways to accommodate the unprecedented demand for compute and storage resources, including exabyte-level storage infrastructures. As data continues to grow, hyperscale and cloud suppliers are integrating tape storage devices into cloud infrastructure to achieve cost containment and address enormous energy concerns. Just in time, the successful union of archiving, object storage, unstructured data and modern tape has arrived.



Object Storage and Tape Combine to Address Archival Demand

Object storage evolved out of the need to optimize storage performance and scaling capabilities for huge volumes of unstructured data and is the same technology that enables the public cloud. Object storage with its associated metadata can be accessed using HTTP, HTTPS and APIs. For these reasons, object storage has become the de facto standard for storing information in the cloud as object repositories can scale to hundreds of petabytes in a single namespace without much performance degradation. In recent years object storage for tape has been introduced, bringing the best of both worlds together. Simple, long-term, data storage management with low-cost, long-lived tape automation allows for objects and metadata to be efficiently read and written on tape in its native form. Cloud providers are using object storage solutions for their archive services storing large amounts of archival data, cost-efficiently on tape. Object storage momentum is building and [AMA research](#) projects the object storage market to grow at 14.5% CAGR.

LTFS (Linear Tape File System)

LTFS (Linear Tape File System) provides a significant step forward in moving tape storage away from its reputation as complex and difficult to use. LTFS provides an open file system format whereby a user can access files directly without the application that wrote the data. LTFS provides the back-end connector for SwiftHLM (Swift High Latency Media) enabling users to perform bulk tape operations within a Swift data ring. LTFS makes archiving and retrieving object data easier for tape applications and continues to gain momentum as 38 companies are LTFS implementers.



Energy Savings Favor Tape

Data centers and information technology contribute nearly 2% to global carbon emissions and currently consumes over 2% of the world's electricity and is expected to [soar up to 8%](#) by 2030 as concerns about the availability of sufficient power supplies grow. Hyperscale data centers are directly confronting this challenge as the insatiable growth of servers and disk farms are devouring budgets, overcrowding data centers and creating enormous energy and carbon footprint problems. Shifting less active and archival data from disk to tape and virtualizing servers are the two most significant ways of reducing data center energy consumption. Tape cartridges spend most of their life in a library slot or on a shelf and don't consume energy unless mounted in a tape drive, making tape the ideal archival storage choice. Building another data center is extremely expensive mandating that energy consumption be efficiently managed accentuating the significant role tape is playing in data center economics. For large-scale data centers adding disk is tactical – adding tape is strategic.

Tape TCO (Total Cost of Ownership) Calculators Become Available

Tape's cost per terabyte and TCO advantage compared with other storage mediums makes it the most cost-effective technology for long-term data retention. Easy to use and publicly available TCO calculators are available from [Fujifilm](#) and the LTO consortium. These tools allow users to define input assumptions to help assess the TCO of automated tape systems compared to HDDs and cloud storage. See example TCO comparison below from the [LTO.org TCO tool](#).

SOLUTION 10 YEAR TCO

TAPE WITH SDS SOFTWARE	\$1,994,512.99
ON-PREMISE OBSTOR JBOD	\$13,608,102.55
HYPERSCALE	\$6,043,464.73
ORACLE ARCHIVE	\$2,986,763.94
WASABI	\$13,219,520.18

	Tape	On-premise JBOD	Hyperscale	Oracle	Wasabi
Storage Expense	\$736,614.40	\$4,920,659.58	\$5,490,394.80	\$2,436,471.57	\$12,669,520.18
Capital Costs	\$656,376.19	\$6,522,734.79	\$0.00	\$0.00	\$0.00
Administration/Maintenance Costs	\$500,000.00	\$1,827,334.64	\$550,000.00	\$550,000.00	\$550,000.00
Command Charges	\$0.00	\$0.00	\$4,385.05	\$417.62	\$0.00
Egress charges	\$0.00	\$0.00	\$0.00	\$0.00	\$0.00
Energy	\$60,000.00	\$292,373.54	\$0.00	\$0.00	\$0.00
Floor Space	\$41,522.40	\$45,000.00	\$0.00	\$0.00	\$0.00

Tape Systems Improving Access Time and Throughput

In addition to tape’s continual capacity improvements, tape has improved file access times (time to first byte) and data rate (throughput) with Active Archive, LTFS, RAO, TAOS, RAIT and RAIL, while offering the storage industry’s fastest data rates. Initial file access time reductions of ~50% are typical.


Tape Performance Improvement Summary	
Access Time Improvements (Time to 1 st Byte)	
Active Archive	Active Archive is gaining momentum and greatly improves access time to tape data by using high capacity HDDs or SSDs as a cache buffer in conjunction with a tape library.
RAO	RAO (Recommended Access Order) for enterprise tape produces an optimized list called “best access order” and provides the least amount of time that is needed to locate and read all files or data sets on a tape.
TAOS	TAOS™ (Time-based Access Ordering System) the order in which the files are retrieved can be reordered for “best access” using shortcuts on the tape reducing the tape movement time required to retrieve a file.
LTFS	Provides access to files directly without the application that wrote the data.
Faster Robotics	Intelligent robotics optimize robotic movements improving access time and reliability.
Throughput and Availability Improvements	
RAIT	RAIT (Redundant Arrays of Independent Tape) aggregates bandwidth across multiple tape drives significantly increasing throughput and availability.
RAIL	RAIL (Redundant Arrays of Independent Libraries) aggregates bandwidth across multiple libraries increasing throughput and providing geographic resiliency.

Tape Addresses Storage-Intensive Applications and Workflows

In addition to tape’s continual capacity improvements, tape has improved file access times (time to first byte) and data rate. Tape has become the optimal storage solution for many next-generation applications that are quickly exceeding the capabilities of traditional infrastructures. This includes addressing the storage and data security requirements for Big Data, cloud storage services, entertainment, hyperscale computing, IoT, and surveillance that are all projected to drive enormous high-value storage demand. Estimates suggest as much as 90 percent of data created is rarely touched once it’s been stored. Key workflows and applications requiring long-term storage impacting tape growth include:



Big Data encompasses many disciplines and is the massive amount of data that inundate organizations on a constant basis. The [next-generation data storage market](#) is estimated to grow from \$53.9 billion in 2020 to \$81.0 billion by 2025 at a CAGR of 8.5%.



The [global cloud storage](#) market size is projected to grow from \$50.1 billion in 2020 to \$137.3 billion by 2025, at a CAGR of 22.3% during the forecast period. The cloud storage market is expected to rise mainly due to growing data volumes across enterprises, the rising need to provide the remote workforce with ubiquitous access to any data type, and the potential TCO and cost-savings of cloud storage solutions.



Online banking transaction archives, ATM history, POS, audit and communication logs, and compliance regulations generate data with long-term retention requirements.



Electronic medical records, images (X-Ray, MRI, CT), genome sequences, pharmaceutical development and approval data, and telemedicine drive many new use cases. LTO with LTFS allows clinicians and administrators to quickly retrieve and share EMR, PACS, DICOM and other medical data.




High Performance Computing uses archives to feed compute-intensive applications for pattern recognition and simulate future consequences and predict outcomes. When the study is complete, the data becomes archival again awaiting future analysis. The [total HPC storage market](#) worldwide was \$5.8 billion in 2019, up 2.7% from 2018.



Long-term accident records, images, claims, disputes, and payment history are kept indefinitely.



Total IoT spending neared the [\\$745 billion mark in 2019](#) and is now viewed as a main driver of the digital transformation. The IoT, mobile apps, autonomous vehicles, video, RADAR, LIDAR and sensor data will generate data much faster than it can be analyzed creating an enormous archive pile up for future analysis. By 2022, the IoT is expected to surpass \$1 trillion from new electronic products creating many new storage and security requirements for data to be analyzed at a later time.



The M&E industry relies heavily on tape and digital archives to protect raw production footage. For movie production, it is common to have workflows that need access to archived digital assets at some point in time after the movie is made. Most M&E content is never deleted and uses and repurposes archival content to reach new customers and create new revenue streams. The M&E industry has extensive experience with film archiving and data migration as preserving digital content is a mission-critical function for their survival. [Over 174 exabytes](#) of new digital storage are projected to be used for M&E digital archiving, content conversion and preservation by 2024.



Raw camera footage typically becomes archival after 7 days and surveillance retention periods are quickly increasing. The [airborne surveillance](#) market is projected to reach \$5.8 billion in 2023. The [average surveillance data](#) generated daily by video cameras globally, before cutting and editing, is estimated to reach 3.5 exabytes in 2023.



Archives provide research input and potentially new results, including data for seismic tests for oil & gas exploration, atmospheric science and predictive weather modelling.



The MLB Network alone archives over 1.2 million hours of content, which is indexed and stored with infinite retention periods and makes it available to the production team via proxy video.

Tape Value Proposition is Compelling

Continued development and investment in tape library, drive, media and data management software has effectively addressed the relentless demand for improved reliability, higher capacity, better power efficiency, ease of use and the lowest cost per GB and TCO of any storage solution. Below is a summary of tape’s value proposition followed by key metrics for each function.

Function	Current State of the Tape Industry
Price/TCO	Tape has the lowest acquisition price \$/GB, HDD TCO 5-8x higher than tape.
Performance & Access Time	Much Improved - Active Archive, fastest data rates, RAIT, smarter and faster robotics, access time to 1st byte features (RAO, TAOS) have arrived. Few HDD performance gains identified.
Recovery Speed	Faster recovery for large data sets and files than HDD.
Capacity	Tape cartridge capacity max. @20TB (60TB compressed) with 200X more surface area than HDD @14TB. Lab demos demonstrate native cartridge capabilities to reach 100s of TB.
Data Density (Floorspace)	The largest capacity robotic library stores an exabyte (1x10 ¹⁸) in 2.558 PB per sq. ft. (391 sq. ft.) with 55,900 LTO-9 18 TB cartridges.
Scalability	Tape adds capacity by adding media, HDDs add capacity by adding drives.
Reliability	Tape (BER) 1x10 ²⁰ HDD (BER) 1x10 ¹⁶ Tape BER is 10,000x higher than the best HDD.
Data Integrity	Tape uses a read-while-writing process that checks data as it is written on tape. HDDs provides full verify essentially resulting in a 50% data rate reduction.
Energy Usage	Tape uses much less energy than HDDs, can move tape data without electricity.
Portability	Tape media easily portable during electrical outage, HDDs more difficult to move.
Encryption/WORM	All tape drives offer Encryption/WORM, optional on some HDDs but seldom used.
Cybersecurity	Tape air gap prevents cybercrime attacks, strong defense against malware.
Media Life	>30 Years for all modern tape, average ~4-5 years for HDDs.
Recording Limits	No foreseen limits for tape, HDDs facing areal density and performance limits.
Cloud Storage & Hyperscale	Tape improves cloud reliability and security, lowers storage costs, unlimited capacity scaling. Hyperscale will require tape to manage extreme growth.



Using Tape for Cybersecurity Prevention

Security and cybercrime protection have gained serious attention in recent months. [Ransomware attacks](#) are on the rise and they're estimated to cost global organizations \$20 billion by 2021 with government agencies, healthcare providers, and educational institutions in the U.S. impacted by ransomware attacks at a cost of more than \$7.5 billion in 2019 alone. As ransomware attacks become more targeted and damaging, your organization faces increased risk that can have your networks down for days or even weeks.



The Tape Air Gap Provides Data Security and Cybercrime Prevention

The tape air gap, inherent with tape technology, has ignited and renewed interest in cybercrime prevention. The “[tape air gap](#)” means that there is no electronic connection to the data stored on the removeable tape cartridge therefore preventing a malware attack on stored data. HDD and SSD systems remaining online 7x24x365 are always vulnerable to a cybercrime attack. [Cybercrime](#) is expected to cost the world \$6 trillion in 2021 with ransomware attacks occurring every 11 seconds. Air gapping should be an integral part of any archive, backup, recovery and security plan. In addition to cybercrime protection, tape provides a secure offline solution for backups. The best practice data protection scheme is often described as the 3-2-1 rule: at least 3 copies of the data, on at least 2 different media technologies (i.e.: one on HDD, one on tape), and at least 1 copy that is offline and offsite.

Additional Resources



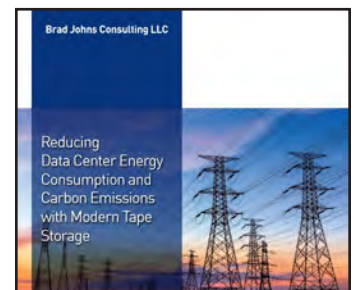
[How To Store An Exabyte Spectra e-book](#)



[Quantum Ransomware Solution Brief](#)



[Horison The Tape Air Gap White Paper](#)



[Reducing Data Center Energy Consumption with Tape](#)



[2019 Tape Roadmap INSIC Report](#)



[Horison The Ascent to Hyperscale White Paper](#)



[Leveraging Tape to Combat Ransomware Report](#)



[Horison Tape Storage White Paper](#)



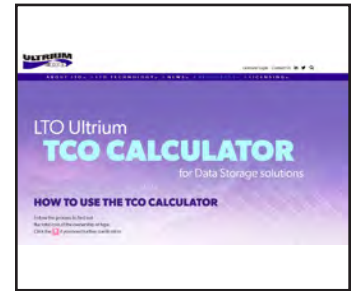
[Horison Tiered Storage White Paper](#)



[Horison Tape Performance Accelerates White Paper](#)



[FUJIFILM TCO Tool](#)



[LTO TCO Calculator](#)



[Active Archive 2020 Annual Report](#)



[PoINT Max Planck Institute Bad Nauheim White Paper](#)



[ESG LTO White Paper](#)



[Spectra Data Storage Outlook Report](#)

Summary

Even with a global pandemic in full swing and how it plays out remaining unclear, the digital data storage landscape continues to expand presenting a constant stream of new challenges. This is evident for the tape industry as the steady arrival of many rich technology improvements has set the stage for tape to remain the most cost-effective storage solution for the enormous high capacity and archival challenges that lie ahead. Today's data center storage technology hierarchy consists of three technologies - SSDs, HDDs and tape and the ideal storage implementation optimizes the strengths of each. However, the role tape serves in today's modern data centers is expanding fast and tape momentum will increase as data growth continues on an explosive trajectory across many new applications, workloads, and in most of the hyperscale and hyperscale-lite data centers. With these advancements in place, modern tape technology delivers the most reliable, energy efficient and cost-effective data center storage solution available today. Roadmaps signal this trend of steady technological tape innovation to continue well into the future.