

## Tape to Play Critical Roles as the Zettabyte Era Takes Off

Tape Advancements Push Storage and Sustainability Benefits to New Levels



### May 2022

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, MagStor, mLogic, Park Place Technologies, Oracle, Overland-Tandberg, OWC Mercury, PoINT Software & Systems, Qualstar, Quantum, Spectra Logic, StrongBox Data Solutions, StorageDNA, SullivanStrickler, Symply, Turtle, and XenData, has issued this report to highlight the current trends, usages and technology innovations occurring within the tape storage industry.

### Introduction


The zettabyte era is in full swing generating unprecedented capacity demand as many businesses move closer to exascale storage requirements. The ever-increasing thirst for IT services has pushed energy usage, carbon emissions and keeping the storage industry's growing impact on global climate change in check to center stage. The modern data center operates with little human activity but consumes at least [1% of the world's power production](#) and will only continue to grow with new applications. The compute-intensive process of mining [Bitcoin](#) alone uses around 91 terawatt-hours of electricity annually, more than Finland, a nation of about 5.5 million people. The impact from global data centers on climate change has become a central challenge for the entire IT industry. Looking ahead, global climate change and carbon emissions will begin to eclipse the Covid pandemic and the growing value of eco-friendly tape will become even more compelling. In addition to reducing carbon emissions, tape plays a central role in lowering data center TCO and containing storage growth while offering an additional level of cybercrime protection. Expect tape to play an even broader role in the entire IT ecosystem going forward as the zettabyte era takes off.



# Managing the Zettabyte Era Data Growth

Data is being generated faster than it can be analyzed, significantly extending data retention timeframes. This amount of data quickly becomes cost-prohibitive to maintain online all the time, since 60 - 80% of all data stored is infrequently accessed and classified as archival or cold data. As storage requirements soar, the energy and economic value of moving cold data from energy-intensive HDDs to tape becomes obvious. IDC's Global Data Stored chart projects 11.1 ZBs of data will be stored in 2025, suggesting between ~7.0 - 9.0 ZBs of all data stored in 2025 would be classified as archival or cold, and most cost-effectively stored on tape. Many large organizations are seeing their stored data double every two years. One ZB (zettabyte =  $1 \times 10^{18}$ ) would fill 55.56 million LTO-9 (18 TB) cartridges or 18 TB HDDs, or 125 million years of one-hour TV shows.

Global Data Stored, 2015-2025



Source: IDC White Paper, sponsored by Fujifilm, September 2021  
Accelerating Green Datacenter Progress with Sustainable Storage Strategies, doc #US48252321

## LTO Tape Supports Zettabyte Era Growth with Record Capacity Shipments in 2021

The LTO Program Technology Provider Companies (TPCs), Hewlett Packard Enterprise Company, IBM Corporation and Quantum Corporation, released their annual tape media shipment report, detailing year-over-year shipments through the fourth quarter of 2021. The report reveals 148 Exabytes (EB) of total tape capacity (compressed) shipped in 2021, marking an impressive record year for the LTO Program. With a growth rate of 40%, this strong performance in shipments continues following the previous record-breaking 114 EB capacity shipped in 2019 and 105 EB of capacity shipped in the pandemic affected year of 2020. Organizations continue to invest in LTO tape technology thanks to its high capacity, reliability, low cost, low power consumption and strong data protection features, especially as threats to cybersecurity soar.

TOTAL CAPACITY BY CY\*\* (EB COMPRESSED)




\* Aggregate capacities do not include LTO-7 Type M media  
\*\* Graph shows data from past 10 years only

## LTO-9 and Higher Performance Features Arrive

Tape technology highlights include the availability of higher capacity LTO-9 tape drives, access time and data rate improvements, widespread support for S3 tape as object storage targets, and increased energy efficiency advantages compared to other storage solutions. The availability of LTO-9 provides a significant boost for cost-effective, scalable, secure storage to address many rapidly evolving 21st century storage-intensive applications. LTO-9 increases the native cartridge

capacity of LTO-8 by 50% to 18 TB (45 TB compressed) and increases drive throughput (11%) up to 400 MB/sec enabling a single LTO-9 drive to write up to 1.44 TB/hour. A new feature with LTO-9, oRAO (Open Recommended Access Order) reduces initial file access times to first byte of data. New R/W head and servo technologies that record even narrower data tracks and further increase cartridge capacity have also been developed. A native 18 TB LTO-9 cartridge can hold about 21 months of data from the Hubble Space Telescope, 9,000 hours of movies or 6,480,000 songs. The current LTO roadmap has considerable headroom and projects a native cartridge capacity up to 144 TBs for LTO-12. The capacity increases are welcome for enterprise, cloud, HPC and hyperscale data center growth as archival data continues to pile up much faster than it is being analyzed.

## LTO ULTRIUM ROADMAP ADDRESSING YOUR STORAGE NEEDS




**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.

## Sustainability - Tape Should Be Included in Every Green Data Center Strategy

As the world faces a growing climate crisis, many data centers have established aggressive sustainability goals to reduce energy consumption and CO2 emissions. Servers and HDDs are the two largest consumers of data center energy. In major data centers, energy consumption increased by 31% from 2017 to 2020, and the amount of data stored is expected to grow by 27% each year through 2025. IDC estimates that globally, migrating more archival and cold data from HDDs to tape can reduce CO2 emissions by 43.7%, or 664 million tons, by 2030.


## CO2 Emissions Avoided by Using Tape Storage



Source: IDC White Paper, sponsored by Fujifilm, September 2021  
 Accelerating Green Datacenter Progress with Sustainable Storage Strategies, doc #US48252321.

## TCO, Energy and Carbon Emissions Reduced With Tape

Carbon Emissions Reduction – A research paper issued by Brad John’s Consulting focused on the entire product lifecycle, entitled, “Improving Information Technology Sustainability with Modern Tape Storage.” This report compared an all data on HDD solution to an all tape solution and an active archive that moved 60% of the HDD resident data to tape. Carbon emissions were reduced by 57% and electronic waste was reduced by 48%. In addition, if all the data was archival/cold and is transferred to tape, carbon emissions are reduced by 95% and TCO lowered by 73%. Tape offers considerable sustainability and cost advantages while addressing the growing impact that data centers are adding to climate change. LTO.org offers an easy-to-use TCO calculator to provide a customized data storage cost profile.



Source: Brad Johns Consulting, LLC, July 2021 “Improving Information Technology Sustainability with Modern Tape Storage”

# Tape Hardens Ransomware and Cybercrime Defense

Ransomware has become a nightmare for everyone and no computer system is immune from cybercrime. There is a new attack somewhere on the Web every 39 seconds, or approximately 2,244 daily attacks on the internet. In 2020, the average data breach cost was \$3.86 million; however, the highest ever ransomware demand grew to [\\$70 million](#). [Quest Strategic Systems](#) reports the average downtime due to a ransomware attack is 21 days. The number of data breaches through September 30, 2021, has exceeded the total number of events in full-year 2020 by 17 percent (1,291 breaches in 2021 compared to 1,108 breaches in 2020). Ransomware is projected to cost organizations globally up to \$20 billion by year-end 2021 and to exceed [\\$265 billion](#) by 2031.

Tape’s inherent air gap provides additional levels of cybercrime defense. The “[tape air gap](#)” means there is no electronic connection to the data stored on the removeable tape cartridge therefore preventing a direct malware attack (hacking) on tape data. HDD and SSD systems connected and online 7x24x365 are always vulnerable to a cybercrime attack making tape the only viable data center air-gapped storage solution available. Backup was the original data protection strategy, but having one backup copy is no longer sufficient. The popular and genetically diverse 3-2-1-1 Data Protection Strategy states that enterprises should have three copies of backup data on two different media types, one copy offsite and one air gap copy. Combining the air gap copy with encryption and WORM, tape plays an integral role to build and strengthen any data center cyber resiliency strategy.

Performance features such as oRAO and TAOS can improve cyber resilience response times by shortening the time needed to access and recover backup and archive data from tape. Modern tape provides immutable air gapping at scale at a favorable price point for backup and archive data and should be an integral part of any data protection strategy whether on-premises, in the cloud, or hybrid.

The 3 – 2 – 1 – 1 Data Protection Strategy			
3	2	1	1
Copies of Data	Different Types of Media	Copy Offsite	Air Gap Copy
Primary copy Backup copy 2nd Backup copy	(SSD, HDD, Tape)	Cloud, remote vault (SSD, HDD, Tape)	(Tape)

Several tape library vendors are taking air gap protection to the next level by offering libraries with managed partitions consisting of dedicated slots which are invisible to external applications. These isolated partitions contain no drives creating a secure storage location within the library providing an additional barrier to access. They are solely configured by the tape library administrator who may create, modify, delete, or reconfigure partitions to meet any size required.

Since partitioned tapes remain in the library, media handling is avoided. Partitions can only be accessed by the library administrator and are not accidentally accessible to any applications.


## Beyond the LTO Roadmap

The tape industry continues to push data rate, reliability, and media life to record levels surpassing disk drives. In Dec. 2020, IBM and Fujifilm demonstrated a record areal density of 317 gb/in<sup>2</sup> on linear magnetic tape which would yield a native cartridge capacity of 580 TBs using a new magnetic particle called Strontium Ferrite (SrFe). SrFe has improved magnetic characteristics beyond Barium Ferrite (BaFe) and delivers even better signal strength. Like BaFe, SrFe is a chemically stable oxide, indicating it will perform well for long-term data preservation. Laboratory demonstrations suggest relatively few tape technology limitations for the coming decade or more, and may signal the way for a new, very low-cost, deep archive storage tier by 2025.

## Will A New Storage Tier Emerge?

Most of the world’s digital data is stored on SSD, HDD, and tape media. As we approach 2025, momentum is building for the emergence of a new deep archive (very cold) secondary storage tier. Deep archives must be easily scalable and require minimal remastering (long media life cycles) and may be the home for the “golden copy”, an immutable master copy of the most critical data. Remastering is a labor-intensive process that typically migrates data to new media every ~5 - 10 years depending on the technology. Access times of hours or even a few days can be acceptable. The optimal design will have a robust roadmap,

**Metal Particle vs. Barium Ferrite vs. Strontium Ferrite**



minimal environmental/carbon footprint, and a price point below all other competitive offerings. Deep archives will differ from traditional archival data as it will theoretically contain data that is classified as WORSE (Write-Once-Read- Seldom-if-Ever) or WORN (Write-Once-Read-Never).


The modern tape roadmap has relatively few areal density limitations and is currently the best positioned technology to address a new storage tier. Other technologies are focusing on the deep archive tier and remain under various stages of development including [DNA](#) which requires the process of encoding and decoding binary data to and from synthesized strands of DNA. Though the cost of synthesis, access times and reliability compared to magnetic storage present risks and remain unknown.

## Broader Roles Arrive for Tape - Addressing The Looming Storage Optimization Challenge

Effectively addressing the storage optimization challenge of “getting the right data, in the right place, at the right time and at the right cost” presents the biggest returns for storage managers. To meet this challenge, the strengths of tiered storage with SSDs, HDDs, tape, and cloud must be leveraged. As most data ages, access frequency drops off rapidly and data typically reaches archival status in 90 - 120 days eventually becoming cold data. However much archival data continues to live on HDDs long after it reaches archival status, an expensive residence for cold data. Fortunately, modern tape technology has improved access times, data rates, and reduced TCO and energy consumption making it best suited to address the growing amount of cold data. Intelligent data management software and an active archive are key ingredients enabling an optimized and scalable storage infrastructure to become a reality.

## The Active Archive - Getting the Right Data, In the Right Place, At the Right Time, At the Right Cost

Storage optimization receives a big boost from an active archive which provides dynamic optimization and fast data access for archival storage systems. An active archive integrates two or more storage technologies (SSD, HDD, tape, and cloud storage) behind a file system providing a seamless means to manage archive data in a single virtualized storage pool. SSDs or HDDs serve as a cache buffer for archival data stored on tape providing faster access to first byte of data, higher IOPs, and random access for more active archives. Many data management products now support tape as an object storage target using S3 services. Combining the open tape file system [LTFS](#) with tape partitioning, data mover software (HSM, etc.), and an SSD, HDD array or NAS in front of a tape library creates an active archive. LTFS currently has 36 implementers and it is expected that an increasing number of ISVs (Independent Software Vendors) will exploit LTFS in the future. The active archive concept is supported by the [Active Archive Alliance](#). See the Active Archive conceptual view in the image below.





## Data-Intensive Applications and Workflows Fuel New Tape Growth

The number of vertical markets and more complex use cases effectively addressed by tape continued to expand in 2021 as access time improvements, TCO benefits and significantly less environmental impact combined to deliver a highly effective and secure storage infrastructure. Tape use cases have been shifting from smaller files to larger, linear data streams that require intelligent data management services throughout the data lifecycle. The use cases listed below are thriving in the zettabyte era generating enormous amounts of high-value and unstructured data, much of which is stored as archives awaiting future reference. Given that workflow pattern, many of these applications can immediately benefit from an active archive implementation leveraging high capacity, low cost tape systems.



### Big Data

The global big data [market size](#) is projected to reach as much as \$116 billion by the end of 2027. The increasing integration of AI and machine learning analytics will emerge to accelerate market growth. The market was worth \$41.33 billion in 2019 and is projected to exhibit a CAGR of 14.0% during the forecast period, 2020-2027. Big Data analytics activate cold archives and can realize significant benefits from an active archive.



### Cloud Storage Services

The [cloud storage market](#) size was valued at \$46.12 billion in 2019, and is projected to reach \$222.25 billion by 2027, growing at a CAGR of 21.9% from 2020 to 2027. Growth is expected to rise from a significant shift to hybrid clouds and data lakes as the primary deployment model. Increasing cloud adoption across industry verticals, security concerns over cloud storage, high egress fees, and cloud adoption rates among small and medium enterprises are all significant factors that impact growth of the global cloud storage market. It's estimated that [82% of enterprises](#) today use hybrid (on premises and public) clouds.




### HPC

High Performance Computing often uses active archives to feed compute-intensive applications for analytics, pattern recognition, and simulations to predict outcomes. Despite the turmoil and impact resulting from the pandemic, the broad [HPC market](#) finished 2020 at \$38.9 billion in revenue, down just 0.2 percent from 2019. HPC is returning to growth and is projected to reach \$60 billion in 2025.



The total [number](#) of IoT connected devices is expected to reach 10.07 billion by the year 2021 and reach 25.44 billion by year end 2030. Spending neared the \$761 billion mark in 2020 and is now viewed as a main driver of the digital transformation. The IoT, mobile apps, autonomous vehicles, video, LIDAR, RADAR, drones and sensors will generate data much faster than it can be analyzed leaving an enormous archive data pile up ideally stored on tape for future analysis. The [IoT market](#) expected to reach \$1.3 trillion creating countless new storage and data security requirements.



The life science [analytics market](#) size was valued at \$7.7 billion in 2020 and forecast to reach \$14.1 billion by 2028, a 7.8% CAGR. Increasing use of big data analytics and predictive analysis to identify future trends and outcomes are keys for growth. Medical records, images (X-ray, MRI, CT), genome sequences, pharmaceutical development and approval, telemedicine and Covid vaccine research are driving new use cases. LTO with LTFS allows clinicians and administrators to quickly retrieve and share EMR, PACS, DICOM and medical data that often resides in data lakes for years before reference.



The [media and entertainment industry](#), a major tape market, was impacted throughout 2020 with many live events canceled because of the Covid-19 pandemic. For movie production, it is common to have workflows that rely heavily on secure tape to protect and permanently archive raw production footage. The global media and entertainment storage market size was valued at \$6.6 billion in 2019, and is projected to reach \$17.9 billion by 2027, a CAGR of 13.7%. More than [174 exabytes](#) of storage capacity is projected for M&E digital archiving, content conversion and preservation by 2024.



Video surveillance is now the largest video market, and it has become evident everywhere. The proliferation of cameras and surveillance systems across all verticals promises significant and sustainable growth for the video surveillance storage market. More camera counts, higher resolutions, AI, and metadata-based analytics are straining security infrastructures, driving needs for more scale, more performance, and the latest storage technologies. Raw camera footage typically becomes archival after 7 days and surveillance retention periods are quickly increasing. The global daily average surveillance data generated, before cutting and editing, is estimated to reach 3.5 exabytes in 2023. The [video surveillance](#) market is projected to grow from \$7.5 billion in 2020 to \$10.2 billion by 2025, a CAGR of 6.1%.



The MLB Network alone archives over 1.2 million hours of video content, which is indexed and stored with infinite retention periods and makes it available to the production team via proxy video. A [four-hour MLB game](#) typically consumes between 1.43 GB and 5.45 GB depending on resolution. In addition to storing and subsequently archiving most sporting event broadcasts forever, the global sports analytics market size is expected to grow from \$1.9 billion in 2019 to \$5.2 billion by 2024, at a CAGR of 22.0%. Major growth drivers of the sports market come from professional and collegiate teams which are increasing investments for data-driven decisions with the goal of improving player and team performance.



## Tape Performance Improves Access Times and Throughput

HDDs and SSDs have faster access times to the first byte of data. For large files, tape systems have faster access times to the last byte of data. A new feature available with LTO-9, [oRAO](#) (Open Recommended Access Order) reduces initial file access time serving as a data retrieval accelerator enabling applications to retrieve non-consecutive files from tape by optimizing physical seek times between files. oRAO can reduce initial LTO-9 file access times over 70% while significantly reducing physical tape movement and drive wear by creating an optimally ordered list of files on a cartridge. LTO-9 increases drive data rate to 400 MB/sec. making it ideal for data streaming and large file transfers.

Tape Performance Improvement Summary	
Access Time Improvements (Time to 1st Byte)	
<b>Active Archive</b>	<a href="#">Active Archive</a> greatly improves access time to file and object tape data by using high-capacity HDDs or SSDs as a cache buffer in conjunction with a tape library.
<b>oRAO</b>	<a href="#">oRAO</a> (Open Recommended Access Order) is new with LTO-9 and produces an optimized list called “best access order” enabling applications to retrieve non-consecutive files from tape by minimizing physical seek times between files.
<b>TAOS</b>	<a href="#">TAOS™</a> (Time-based Access Ordering System) the order in which the LTO files are retrieved can be reordered for “best access” using shortcuts on the tape reducing the tape movement time required to retrieve a file.
<b>LTFS</b>	<a href="#">LTFS</a> provides access to files directly without the application that wrote the data.
<b>Faster Library Robotics</b>	Faster, intelligent robotics optimize robotic movements reducing cartridge mount and access times while improving reliability.
<b>Fastest Data Rates</b>	The LTO-9 and TS1160 enterprise drives each have a data transfer rate of 400 MB/sec. This compares to the 7,200 RPM HDDs ranging between 160 – 260 MB/sec.
<b>RAIT</b>	<a href="#">RAIT</a> (Redundant Arrays of Independent Tape) stripes data across multiple tape drives in parallel significantly increasing throughput and provides parity for data reconstruction like RAID does for HDDs.
<b>RAIL</b>	RAIL (Redundant Arrays of Independent Libraries) stripes data across tape cartridges but in different libraries which may be in different geographic locations.

# The Tape Value Proposition

The current state of the tape industry is highlighted with continued development and investment in smart libraries, new drives, advanced media, and intelligent data management software. The tape value proposition below is compelling and addresses the relentless demand for higher reliability, higher capacity, better power efficiency, ease of use and the lowest \$/TB and TCO of any available storage.

Tape Performance Improvement Summary	
Function	Current State of the Tape Industry
<b>Price/TCO/CO2</b>	Tape has the lowest acquisition price \$/TB, lowest TCO and lowest carbon footprint.
<b>Performance</b>	Much Improved – Active archives, fastest data rates, RAIT, smarter and faster robotics, time to 1 <sup>st</sup> Byte features (oRAO, TAOS) can reduce tape access times as much as 70%.
<b>Capacity</b>	LTO-9 cartridge capacity @ 18TB (45TB compressed) with 400MB/sec data rate. Exabyte capacity libraries are available. Lab demos demonstrate tape capacities can reach 580TBs.
<b>Open Standards</b>	LTO tape drives and LTFS provide open standard interface and access.
<b>Scalability</b>	Tape scales capacity by adding media while adding energy consumption. HDDs add capacity by adding drives with adding energy consumption.
<b>Energy/CO2</b>	Tape uses much less energy and has much lower carbon footprint than HDDs (~85%). Playing key role for attaining corporate sustainability goals.
<b>Portability</b>	Tape media easily portable, HDDs difficult to move, in case of data center outages.
<b>Data Integrity</b>	Tape uses a read-while-writing process that checks data as it is written on tape. HDDs provide full verify essentially resulting in a 50% data rate reduction.
<b>Cybersecurity</b>	Tape air gap, encryption, and WORM provide defense against cybercrime and malware.
<b>Durability and Reliability</b>	LTO reliability (BER 1X10 <sup>19</sup> ) has surpassed HDDs (1X10 <sup>16</sup> ), media life > 30 years for modern tape.
<b>Recording Limits</b>	HDDs face areal density and performance limits. Defined and sustainable roadmap for tape.
<b>Cloud Storage</b>	Tape improves cloud reliability and security, lowers storage costs and unlimited capacity scaling.


## Supporting Resources



[Accelerating Green Datacenter Progress with Sustainable Storage Strategies](#)



[Improving Information Technology Sustainability with Modern Tape Storage](#)



[Reducing Data Center Energy Consumption and Carbon Emissions with Modern Tape Storage](#)



[How Tape Technology Delivers Value in Modern Data-driven Businesses](#)



[Saved by the Data. Active Archive Leads the Way in a Mid-Pandemic World](#)



[Protecting Your Business from Ransomware](#)




[Data Storage Outlook](#)




[How to Store an Exabyte](#)



[Hyperscale Storage and Energy Consumption](#)



[Understanding the Data Life Cycle](#)



[Is the IOT, Edge & Fog the Next Frontier for Tape Storage?](#)

## Conclusion

As the zettabyte era accelerates, it's anticipated that the issue of carbon emissions and climate change will become even more critical. Fortunately the tape ecosystem has significantly expanded its roles in recent years evolving from its original function as a backup target to a key storage optimization solution contributing significantly to improving carbon emissions from IT operations. While the sustainability benefits of tape are substantial, the corresponding TCO savings from moving cold data from HDDs to tape are also quite impressive. Tape has become the leading pure storage solution to defend against cybercrime by seamlessly integrating air gap, encryption and WORM capabilities. New capacity, performance and throughput advancements have enabled tape to address many storage-intensive vertical markets that were previously anchored in the domain of costly HDDs. With these advancements in place, modern tape technology delivers the most reliable, energy-efficient and cost-effective storage solutions available today and has the greatest potential to address the massive capacity demand of the zettabyte era. Roadmaps signal that this trend of steady tape innovation will continue in 2022 and well into the future.

[Tape Storage Council Members – current list](#)