Tape Storage

It's a New Game With New Rules

Fred Moore
President
Horison Information Strategies
Horison.com

Abstract

Modern tape storage has become the leading strategic and lowest-cost storage solution for massive amounts of archival and unstructured data. This bodes well for future tape growth as archival data is piling up much faster than it is being analyzed. Over the past decade, the magnetic tape industry has successfully re-architected itself delivering compelling technologies and functionality including cartridge capacity increases, vastly improved bit error rates yielding the highest reliability of any storage device, a media life of 30 years or more, and faster data transfer rates than any previous tape or HDD (Hard Disk Drive). Many of these innovations have resulted from technologies borrowed from the HDD industry and have been used in the development of both LTO (Linear Tape Open) and enterprise tape products. Additional tape functionality including LTFS, RAIT, RAO, TAOS, smart libraries and the Active Archive adds further value to the tape lineup. HDD technology advancement has slowed while progress for tape, SSD (Solid State Disk) and other semiconductor memories is steadily increasing. Fortunately, today's tape technology is nothing like the tape of the past. For tape it's clearly a new game with new rules!

The Tape Renaissance Transitioned to the Era of Modern Tape

Since the first tape drives appeared in the early 1950s, tape has primarily served as a backup and archive device for disk data. Troublesome tape issues of the past including edge damage, stretch, tear, loading problems, and media alignment with older (now obsolete) tape formats such as DAT, DDS, DLT, and 8MM tape were successfully addressed. By 2000, the Legacy Tape Era was ending, and the tape technology *renaissance* was underway as the tape industry was building a new foundation to address many new storage intensive applications. While backup remains an active use case for tape due to its value for fast site restores, future growth opportunities lie with new and emerging markets.

With the internet, hyperscale data centers, cloud, big data, tele-health, compliance, analytics and IoT waves all promising unprecedented data growth, the timing for advanced tape functionality couldn't be better. You may have somehow missed it, but the modern tape era has arrived delivering the following capabilities:

- tape is cheaper ($/TB) to acquire than disk
- tape is less costly to own and operate (lower TCO) than disk by 5-8x
- tape is more reliable than disk by at least three orders of magnitude
- the media life for modern tape is 30 years or more for all new media
- faster tape drive performance (throughput) with RAIT, and access times with RAO and TAOS
- tape libraries are delivering intelligent, faster, and more efficient robotic movement
- LTFS is a standard open tape file system with media partitions for faster "disk-like" access
- the 10-year roadmap for tape technology is well defined with few foreseeable limits

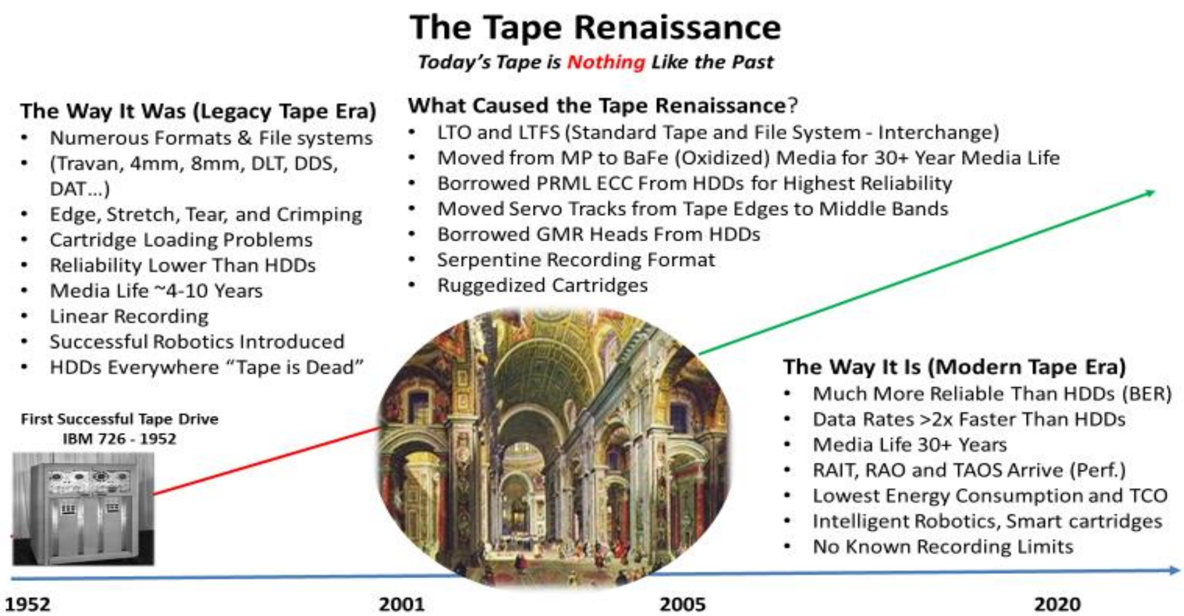

***Key point:* The tape renaissance completely re-architected the mass storage landscape and the modern tape era has arrived. It's time to bring your understanding of modern tape up to date and take advantage of the many benefits that tape has to offer!**

Areal Density is Key for Storage Capacity Growth

Areal density refers to how many bits of information can be stored on a given surface area of a data storage media. To a large degree, the $/GB reduction for HDDs has been driven by the steady areal density growth of hard disk technology for over 30 years. The INSIC chart (Information Storage Industry Consortium) below shows the tape and disk areal density roadmap and demos dating back to 1989 and projects tape growth through to 2029. Between 2003 and 2009, the areal density growth for HDD was approximately 39% per year. More recently, between 2009 and 2018, the average rate of HDD areal density scaling has decreased to 7.6% per year. Looking ahead, there is growing uncertainty regarding future density scaling rates of HDD.

Current LTO and enterprise tape drives operate at areal densities that are about two orders of magnitude lower than the latest HDDs. Therefore, the additional unused tape area makes it possible to advance tape densities at historical rates for at least the next decade.

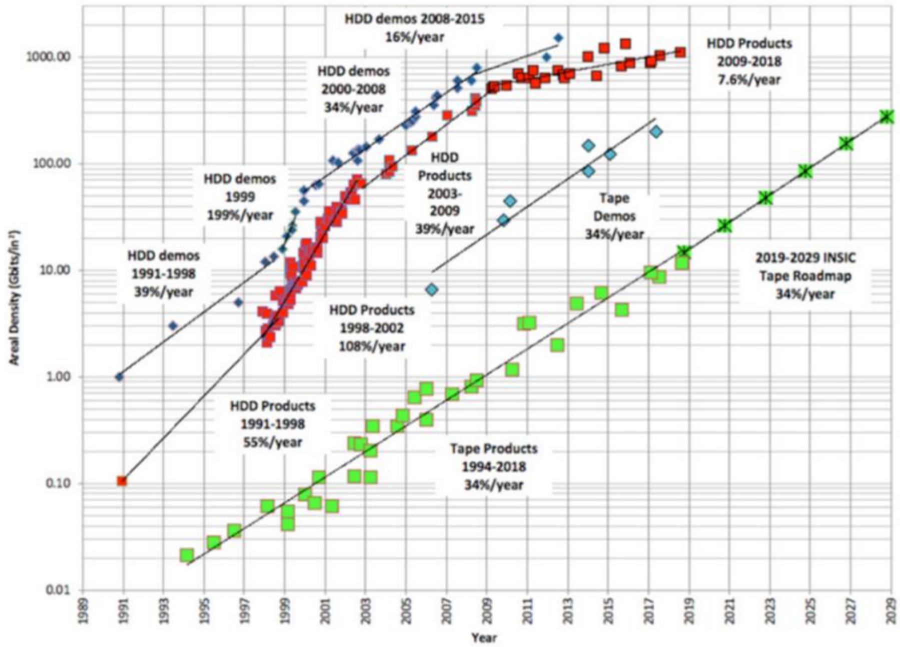

For instance, a 12 TB HDD uses 923 Gb/in^2 vs latest 12 TB LTO-8 cartridge that only uses 8.5 Gb/in^2. Tape can achieve the same capacity with only 1/108th of the areal density than that of the disk leaving room for considerable growth. The latest enterprise TS1160 with TMR (Tunneling Magnetoresistive) heads has a native cartridge capacity of 20 TB and 60 TB compressed (3x), yielding the highest capacity of any storage media. The combination of available tape area and the ability to increase areal densities without major obstacles is the main reason why tape should continue to enjoy the 40% per year capacity growth and 34% per year areal density growth over the next 10 years. With this ability to increase areal density you can expect tape to maintain its cost advantage vs. HDD and other technologies.

Future Data Recording Projections

To increase capacity, some HDDs have increased the number of platters from three to seven while using helium filled disk enclosures to reduce friction. More platters require more heads which means more cost and components that can fail. It should be possible to continue scaling tape areal density at historical rates for at least the next decade before tape begins to face challenges related to the super-paramagnetic effect which today's HDDs have faced. HDDs new energy assisted recording technologies, HAMR (Heat Assisted Magnetic Recording) and MAMR (Microwave Assisted Magnetic Recording) should help further capacity growth. The smaller the magnetic particle, the more data there is in a single bit cell. The net result of these areal density scenarios is a sustained volumetric and native capacity advantage for tape technology.

The future looks much brighter for tape capacity increases. On April 9, 2015 Fujifilm in conjunction with IBM demonstrated a new record in areal density of 123 Gb/in^2 on linear magnetic particulate tape had been achieved. More recently Sony and IBM demonstrated 201 Gb/in^2 with potential for a 330 TB native cartridge. Fujifilm's Strontium Ferrite (SrFe) next-generation magnetic particle recording promises more than 400 TB (33.3 times more storage capacity than LTO-8 at 12 TB) on a cartridge with an areal density of approximately 224 Gb/in^2.

Tape Drive and Media Specifications	Year Introduced	Capacity (native) Compression (x:y)	Data Transfer Rate (native)	Channels per head	Tracks	Areal Density
LTO-6 MP & BaFe	2012	2.5 TB (2:1)	160 MB/sec	16	2,176	2.2 Gb/in^2
LTO-7 BaFe	2015	6.0 TB (2.5:1)	300 MB/sec	32	3,584	4.3 Gb/in^2
LTO-8 BaFe	2019	12.0 TB (2.5:1)	360 MB/sec	32	6,656	8.6 Gb/in^2
TS1140 BaFe	2011	4.0 TB (2.5:1)	250 MB/sec	32	2,560	3.2 Gb/in^2
T10000D BaFe	2006	8.5 TB (2.5:1)	252 MB/sec	32	4,608	4.93 Gb/in^2
TS1150 BaFe	2014	10.0 TB (2.5:1)	360 MB/sec	32	5,120	6.52 Gb/in^2
TS1155 BaFe (TMR)	2017	15.0 TB (2.5:1)	360 MB/sec	32	7,680	9.78 Gb/in^2
TS1160 BaFe (TMR)	2018	20.0 TB (3:1)	400 MB/sec	32	8,704	12.4 Gb/in^2

***Key point:* The tape industry has pushed capacity, data rates, reliability and media life to record levels. Media demonstrations indicate continued advancements in tape technology for many years ahead.**

Tape Leads Storage Reliability Ratings

Customers have indicated for years that a key cause of tape failure was due to media and handling errors; however, these concerns are now out of date. Special prewritten servo tracks allow the tape drive heads to stay aligned with data tracks on the tape to accurately read and write tape data. With the older linear tape products, servo tracks were on the edges of the tape media and dropping a cartridge could often cause servo damage. Since 2000 enterprise and LTO drives have eliminated this issue by combining the pre-recorded servo tracks on the media (between the data bands) along with developing more ruggedized cartridge shells that are relatively impervious to handling damage.

Tape reliability has significantly improved due to several factors. PRML (Partial Response Maximum Likelihood) is the most effective error detection scheme and is widely used in modern disk drives by recovering data from the weak analog read-back signal. LTO drives switched to PRML from the older RLL (Run Length Limited) error checking code. PRML can correctly decode a weaker signal enabling a much higher recording density and allowed tape to surpass disk in reliability. For years MTBF (Mean Time Between Failure) was used to measure storage device reliability but BER (Bit Error Rate) is now the de-facto standard measure of reliability. PRML made it possible for the BER specification for LTO-7 to expect a single undetectable bit error for every 1 x 10^{19} bits transferred. Today, both LTO and enterprise tape products are more reliable than any HDD by three orders of magnitude. Times have changed!

Storage Device Reliability Ratings	BER (Bit Error Rate) Bits read before permanent error
Enterprise Tape (T10000x, TS11xx, LTO-7, 8)	1 x $10E^{19}$ bits
LTO-5-6, SSD (NAND)	1 x $10E^{17}$ bits
Enterprise HDD (FC/SAS)	1 x $10E^{16}$ bits
Enterprise HDD (SATA)	1 x $10E^{15}$ bits
Desktop HDD (SATA)	1 x $10E^{14}$ bits

Source: Vendor's published BER

***Key point:* Tape has the highest reliability of any data storage device and even higher BER levels are expected in the future.**

Tape Roadmap Progress - LTO Roadmap Defined to Generation 12

The [LTO Consortium] publishes a well-defined roadmap (see below) and in Oct. 2017 extended the LTO family to LTO-12. Each successive LTO generation is scheduled to arrive in approximately two- and one-half-year intervals, steadily improving the acquisition price, capacity and performance over previous models. The LTO-7 format expanded the “history buffer” in the compression engine, giving it a 2.5:1 compression ratio, up from 2:1 on previous LTO drives while the enterprise TS1160 drive offers a 3x compression ratio. To push the innovation and capacity boundaries of LTO going forward, the current LTO format required a recording technology transition that supports the higher cartridge capacity growth for future LTO generations. As a result, the LTO-8 specification is only backwards compatible with the former LTO-7 generation rather than two prior versions. A 12 TB LTO-8 cartridge can hold 1,071,000 photos, 12,000 movies or 4,320,000 songs.

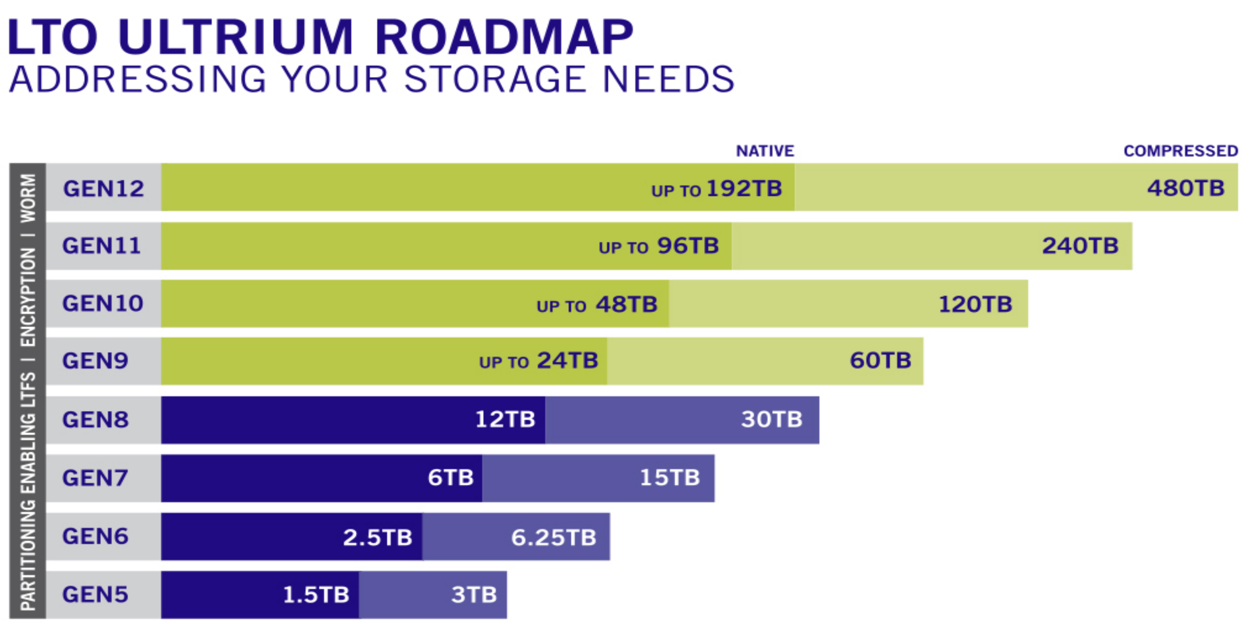

NOTE: Compressed capacity for generation 5 assumes 2:1 compression. Compressed capacities for generations 6-12 assume 2.5:1 compression (achieved with larger compression history buffer).
SOURCE: The LTO Program. The LTO Ultrium roadmap is subject to change without notice and represents goals and objectives only. Linear Tape-Open, LTO, the LTO logo, Ultrium, and the Ultrium logo are registered trademarks of Hewlett Packard Enterprise, IBM and Quantum in the US and other countries.

***Key Point:* The roadmap for future LTO tape systems is well defined, highly attainable, and is expected to support many more years of technology advancements. Expect similar improvements and steady progress for enterprise tape.**

LTFS Enables Faster Data Access for Files, Objects, and Archives

Tape is improving file access times and data rate (throughput) with Active Archive, RAIT, RAO, TAOS and LTFS. LTFS (Linear Tape File System – aka Spectrum Archive) continues to gain momentum and now has 38 companies which are licensed implementers. Developed by IBM and introduced in 2010 with LTO-5, LTFS provides an easier and faster way to access and archive data to tape. LTFS introduced tape partitioning; one partition holds the index and the other contains the content, allowing the tape to be self-describing. The metadata of each cartridge, once mounted, is cached in server memory. Metadata operations such as browsing directory tree structures and searching file names using familiar drag and drop techniques are performed faster in server memory and do not require physical tape movement.

Logical View of LTFS Volume

- **LTFS utilizes media partitioning (LTO-5+ and TS11xx Enterprise).**
- **The LTFS tape is logically divided into partitions "lengthwise".**
- **LTFS places the Index in first partition and Data file(s) in the second partition.**
- **The LTFS index enables faster searching and accessing the files in the second partition via a GUI (Graphical User Interface).**

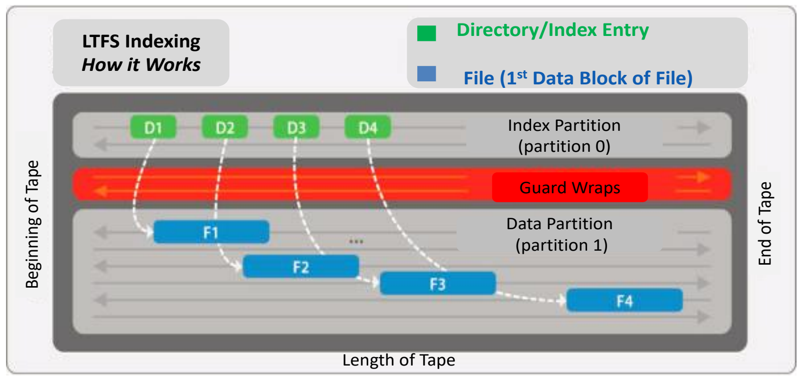

Storing object files on tape has become a reality. Object storage is a common archive and cloud storage format. By bringing object storage to tape, data centers can improve search, scalability, security and lower their TCO beyond that offered by file-based storage solutions while meeting the increasing demand for archival storage. On July 5, 2017 LTFS announced a connection with OpenStack Swift to enable easier movement of cold (archive) object storage data to more economical tape storage for long-term retention. Archive data and cloud storage are major object storage use cases. LTFS now provides a back-end connector for open source SwiftHLM (Swift High Latency Media), a high-latency storage back end that makes it easier for users to perform bulk operations using tape within a Swift data ring. LTFS has made archiving and retrieving object data easier than ever before for tape applications. Object archive software solutions offer a S3-compatible API which is adopted in object storage, so that tape storage can move object files to and from tape libraries in their native object file format and easily integrate with online HDD drive and SSD systems, as well as cloud storage, to optimize the tiered storage model.

***Key point:* LTFS partitioning and its future iterations will provide even greater access capabilities for tape and attract more ISVs (Independent Software Vendors) to exploit its capabilities.**

The Tape Air Gap Provides Security and Cybercrime Prevention

The tape air gap, inherent with tape technology, has ignited and renewed interest in storing data on tape. 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 SDD systems remaining online 7x24x365 are always vulnerable to a cybercrime attack.

Malicious hackers are now attacking computers and networks at a rate of one attack every 39 seconds (University of Maryland). U.S. ransomware attacks cost an estimated $11.5 billion in 2019. The reality is that attackers today have a >90% success rate. Air gapping should be an integral part of any archive, backup and recovery plan. It's not a matter of "if" but "when" hackers will breach your network. There has never been a greater requirement to build cybersecurity data protection than this moment in time. In Sept. 2017, The Wall Street Journal published an article highlighting tape's unique new role in cybersecurity.

Tape Data Rates Accelerate

The INSIC International Magnetic Tape Storage Roadmap projects tape data rates (throughput) to be as much as five times faster than HDDs by 2025. This is great news for businesses and CSPs needing to move, stream or recover large amounts of data the fastest way possible. Data from the big data analytics, higher density images, streaming video, hybrid cloud workloads and traditional DR and hot site applications will benefit from higher data rates. The faster tape data rates will significantly increase with the benefits and value of RAIT, which allows multiple tape drives to transfer data in parallel providing a data rate multiplier. If the fastest data transfer rate (throughput) is needed, tape is the best solution available.

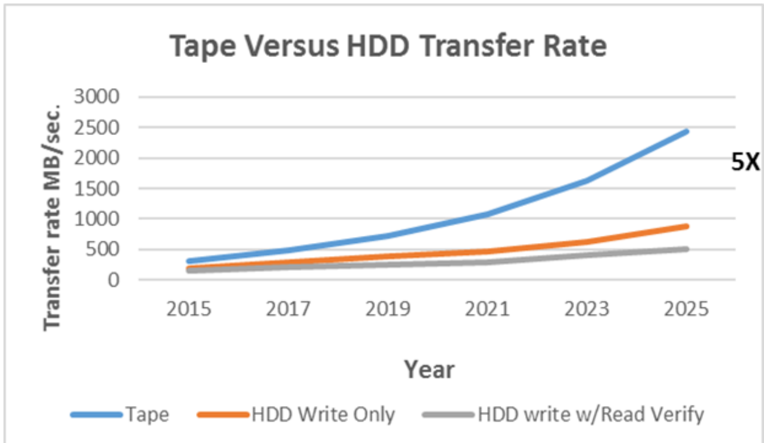

***Key point:* Honestly, did you realize magnetic tape has such a data rate advantage over HDDs?**

RAIT Provides Much Higher Transfer Rates and Tape Access Times Improve

RAIT (Redundant Arrays of Inexpensive Tape) is available with HPSS (High Performance Storage System) and aggregates bandwidth across multiple tape drives in parallel significantly increasing data transfer rate (throughput). RAIT uses multiple tapes loaded in parallel for writing and reading data and provides parity for data reconstruction like RAID does for HDDs. The much higher transfer rates position RAIT for the HPC, HSDC, cloud and enterprise markets.

The tape industry has stepped up its focus to improve tape file access time with two exciting capabilities, RAO (Recommended Access Order) for enterprise tape and TAOS (Time-based Access Order System) for LTO. These features create an optimally ordered list of files on a cartridge which improves file access times as much as 50% while significantly reducing physical tape movement and wear. This capability becomes much more important as tape capacities increase and the probability that the number of concurrently accessed files on a cartridge increases. To complement these features, robotic tape libraries have gotten smarter and faster adding features that minimize robotic time and distance travelled to optimally locate a tape cartridge while adding the benefit of improving library reliability.

Data Center Heat Wave - Tape Means Less Energy Consumption

Computation and data usage have grown dramatically over the last decade leading to an explosion in energy consumption as large data centers strive to keep servers cool and operational. Data centers and information technology consume roughly 2% of the world's electricity currently and is expected to soar up to 8% by 2030. A commonly stated objective for many data center managers today is that *"if data isn't used, it shouldn't consume energy".*

Best practices for using less energy in the data center focus on the two highest areas of energy consumption – servers and disk storage. Tape cartridges spend most of their life in a library slot or on a shelf and consume no energy when not mounted in a tape drive making tape ideal for archival storage. The limits of power consumption in many data centers, especially hyperscale, are being reached forcing organizations to explore new cooling techniques such as water-cooled racks, outdoor and mobile cooling, or in some cases, building another data center. Building another data center is a last resort and is extremely expensive mandating that energy consumption be efficiently managed. Average IT electrical consumption rates for a typical data center is summarized in the chart below. Note: The average US cost of electricity as of December 2019 was $0.1027 per KwH.

Average Electrical Power Usage for Data Centers		
Chillers, cooling, pumps, air-conditioning	**24%**	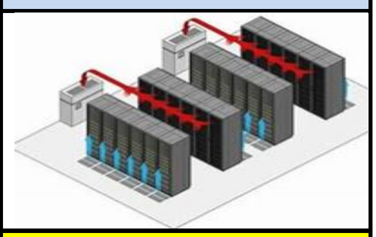
Uninterruptible power supply	**8%**	
Air movement, circulation, fans etc.	**10%**	
Misc. lighting, security, perimeter surveillance	**3%**	
Total infrastructure – external consumption	**45%**	
Servers	**30%**	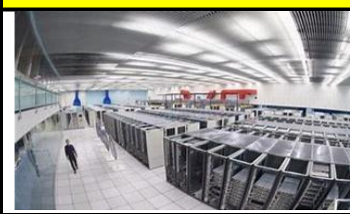
Disk drives, control units	**14%**	
Tape drives, robotic tape libraries	**3%**	
Network gear, SAN switches and other devices…	**8%**	
Total IT equipment – internal consumption	**55%**	

Source: Horison, Inc. and estimates/averages from various industry sources.

Key point:* *Shifting less active and archival data from disk to tape storage and virtualizing servers are the most significant ways of reducing energy consumption in the data center.

Tape Playing an Increasing Role for the Hyperscale and Cloud Rampage

The global HSDC (Hyperscale Data Center) market is forecast to reach $1.47 trillion by 2025 up from $583 billion in 2017. According to the Cisco Global Cloud Index report, the HSDCs are poised to grow from 338 in 2016 to 628 by 2021. Tape's role in the global HSDC and CSP (Cloud Service Provider) market is expanding as tape is quickly becoming the optimal archival storage solution to reduce the much higher HDD operational expense. Using HDDs for archival storage is a strategy – just not a cost-effective one - especially in large scale data centers.

Traditional storage management techniques have left data centers, including hyperscale, struggling with as much as 80% of their data being stored on the wrong tier (HDDs) costing organizations millions of dollars. Few data centers can afford this degree of inefficiency. Hyperscale cloud data centers such as Amazon, Apple, Facebook, Google and Microsoft face enormous energy management challenges which encourage using tape rather than HDDs for storing enormous quantities of less active data. Moving large amounts of archival data in and out of the cloud using expensive network bandwidth can take days or even weeks and can become cost prohibitive compared to moving tape media via truck or airplane. Since tape media is portable, tape becomes highly advantageous if the cloud service provider shuts down or should you want to "quickly" move your entire archive media to another provider.

***Key point:* Hyperscale and cloud providers are deploying tape for their lowest-cost, most secure, long-term archival storage offerings.**

Total Cost of Ownership Favors Tape over Disk

Tape's significant cost per gigabyte and Total Cost of Ownership (TCO) advantage compared with other storage mediums clearly makes it the most cost-effective technology for long-term, secure data retention. Keep in mind that tape capacity can scale without adding more drives – this is not the case with HDDs where each capacity increase requires another drive consuming more energy and cooling For TCO comparisons, the chart below compares automated tape, low-cost disk, and the cloud for tier 3 storage. TCO studies show HDD TCO HDDs typically ranges from 5-8 x higher than equivalent capacity tape systems.

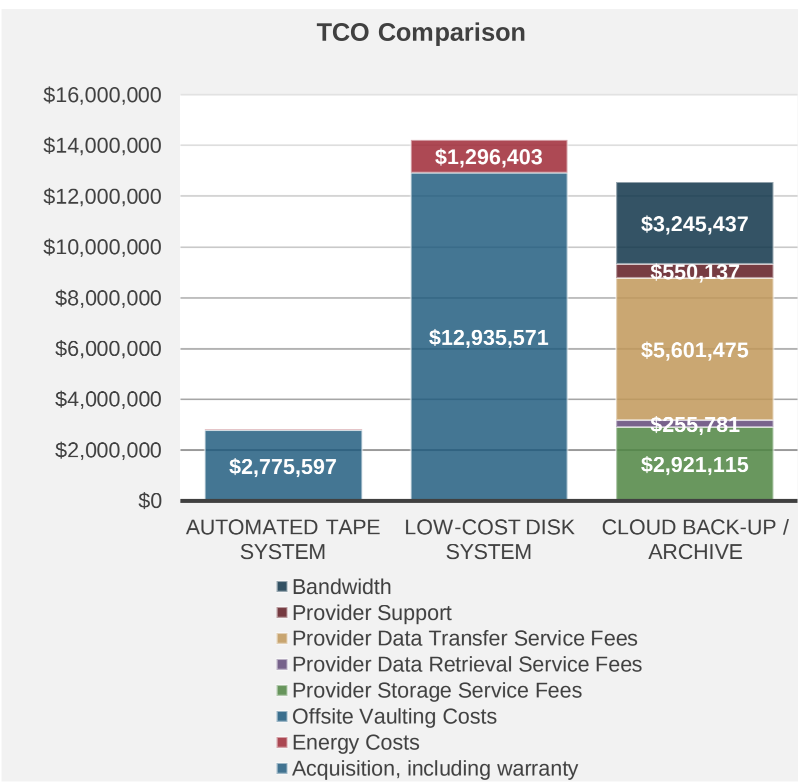

TCO Key Assumptions 20 PB of storage

- 30% CAGR
- 12% of data retrieved year
- LTO 8 Cartridge Price of $112.25
- Technology refreshed in 5 years
- Energy cost of $.105 based on commercial cost per US Energy Information Agency
- TCO includes acquisition, energy, maintenance, cloud storage, network and technology refresh costs
- Tape TCO includes cost of fully automated (robotic) tape library and drive systems

10 Year TCO Results

- Tape Savings versus Disk Storage – 80%
- Tape Savings versus Cloud Storage – 78%

***Key point:* The TCO advantage of tape over disk and the cloud is most compelling for archival storage.**

Comparing Tape and Disk Functionality

Functionality	Tape	Disk
TCO	Favors tape for archive as much as 5-8x over disk	Much higher TCO, more frequent conversions and upgrades
Long-life media	30 years or more on all new enterprise and LTO media favoring archive requirements	~4-5 years for most HDDs before upgrade or replacement, 7-8 years or more is typical for tape drives
Reliability	Tape BER (Bit Error Rate) @ $1x10^{19}$ versus $1x10^{16}$ for disk	Disk BER falling behind - not improving as fast as tape
Inactive data does not consume energy	Yes, this is becoming a goal for most data centers. "If the data isn't being used, it shouldn't consume energy"	Rarely for disk; potentially in the case of "spin-up spin-down" disks *Note: data striping in arrays often negates the spin-down function*
Provide the highest security levels – encryption, WORM	Encryption and WORM available on all LTO and enterprise tape. The tape "air gap" prevents hacking	Becoming available but seldom used on selected disk products, PCs and personal appliances.
Capacity growth rates	Roadmaps favor tape over disk for foreseeable future – native 200+ TB capabilities have been demonstrated	Slowing capacity growth as roadmaps project disk capacity to lag tape for foreseeable future
Scale capacity	Tape can scale by adding cartridges	Disk scales by adding more drives
Data access Time	LTFS, the Active Archive, TAOS and RAO improve tape access time	Disk is much faster (ms) than tape (secs) for initial access
Data transfer rate	400 MB/sec for TS1160, 360 MB/sec for LTO-8, RAIT multiplies data rates	Approx. 160-220 MB/sec for typical HDD
Portability - Move media for DR with or without electricity	Yes, tape media is completely removable and easily transported in absence of data center electricity	Disks are difficult to physically remove and to safely transport
Cloud Storage	Tape improves cloud security, lowers archival storage costs	HDDs become very expensive as cloud & HDSC data centers grow

***Key point:* HDDs are caught in the middle as storage administrators strive to optimize their storage infrastructure to address high performance applications with SSD and archival demands with tape.**

Summary

The past decade has ushered in many key milestones for the tape storage industry. Looking at the decade ahead, expect increasing tape momentum and demand as data growth continues on an explosive path across many applications and workloads, and in most of the major HSDCs. Tape will not replace HDDs or SSDs, but it will be a highly cost-effective complement to SSD and HDD for the foreseeable future. The arrival of many rich tape technology improvements has set the stage for tape to continue to be the most cost-effective storage solution for the enormous high capacity and archival challenges that lie ahead. For tape it's clearly a new game with new rules.