Abstract

The tape industry continues to re-architect itself delivering compelling technologies and functionality including dramatic 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) technology. 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 functionality including LTFS, RAIT, RAO, TAOS and the Active Archive adds further value to the tape lineup. HDD technology advancement has slowed while progress for tape and SSD has been greater over the past 10 years. 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 Era of Modern Tape Arrives Bringing Major Tape Enhancements

Since its inception in the early 1950s, tape primarily served as a backup and archive device for disk data. By 2000, a new era of tape technology was underway as the tape industry was building a foundation to address many new storage intensive applications. 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, Travan, and 8MM tape were successfully addressed. With the internet, cloud, big data and IoT waves 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 delivered the following capabilities:

- tape is cheaper ($/GB) to acquire than disk

- tape is less costly to own and operate (lower TCO) than disk
- 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
- tape drive performance (throughput and access time) improved by RAIT, RAO and TAOS
- tape libraries are delivering faster, more efficient robotic movement
- with LTFS tape has a standard 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: ***News flash – the tape industry has successfully re-architected itself. It's time to bring your understanding of tape up to date and take advantage of the many benefits that tape has to offer*****!**

Reliability Ratings Soar for Tape

Tape reliability has significantly improved due to several factors. LTO drives switched to PRML (Partial Response Maximum Likelihood) from RLL (Run Length Limited). PRML is the most effective error detection scheme and is widely used in modern disk drives. PRML can correctly decode a weaker signal enabling a much higher recording density while allowing tape to surpass disk in reliability.

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 head to stay aligned with data tracks on the tape to accurately read and write tape data. With the older linear tape products, the edges of the tape media served as servo tracks and dropping a cartridge could 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.

For years MTBF (Mean Time Between Failure) was used to measure storage device reliability but this has given way to bit error rate (BER) as the de-facto standard measure of reliability. Starting with LTO-7, several new data-integrity checking features were added that include error-correcting code (ECC) for tape headers. This made it possible for the BER specification for LTO-7 to be increased to no more than 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. 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	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.***

Barium Ferrite Arrives – A Game Changer for Tape Media

For years, Metal Particulate (MP) pigment was the primary tape media type. MP is mainly made of iron (Fe) therefore it will oxidize over time and its magnetic property will deteriorate. BaFe (Barium Ferrite) was developed by Fujifilm providing a major breakthrough for tape media. BaFe is made of an oxide and therefore it does not lose its magnetic property over time due to oxidation. The much smaller BaFe particles are one of the main advantages for using BaFe as it allows more particles-per-unit volume and therefore improves the Signal to Noise Ratio (SNR) and reliability. All generations of LTO cartridges prior to LTO-6 have exclusively used the MP pigment. LTO-6 uses both MP and BaFe while BaFe is now standard on LTO-7 and beyond. BaFe has propelled modern tape media to a life of 30 years or more based on accelerated life tests.

Areal density refers to how many bits of information can be stored on a given surface area of a magnetic disk drive or tape media. On April 9, 2015 Fujifilm Recording Media USA in conjunction with IBM demonstrated (not announced) a new record in areal density of 123 Gb/in^2 on linear magnetic particulate tape had been achieved. This was the fourth time in less than 10 years that Fujifilm and IBM have combined to achieve record-breaking storage capacities on tape. This density breakthrough equates to a standard LTO cartridge storing up to 220 TB of uncompressed data, more than 18 times the capacity of an LTO-8 tape. This tape capacity can preserve the human genome (~600 GB) of 360 people on a single cartridge. The enterprise TS1155 media has a native capacity of 15 TB and 45 TB compressed (3x), yielding the highest capacity of any storage media. A new metal particle is under development by Fujifilm called Strontium Ferrite (SrFe) to ensure steady areal density gains for the foreseeable future.

Tape Drive and Media Specifications		Capacity (native)	Data Transfer Rate (native)	Channels /head	Tracks	Areal Density
LTO-6	MP *and* BaFe	2.5 TB	160 MB/sec	16	2,176	2.2 Gb/in^2
LTO-7	BaFe	6.0 TB	300 MB/sec	32	3,584	4.3 Gb/in^2
LTO-8	BaFe	12.0 TB	360 MB/sec	32	6,656	8.6 Gb/in^2
TS1140	BaFe	4.0 TB	250 MB/sec	32	2,560	3.2 Gb/in^2
T10000D	BaFe	8.5 TB	252 MB/sec	32	4,608	4.93 Gb/in^2
TS1150	BaFe	10.0 TB	360 MB/sec	32	5,120	6.52 Gb/in^2
TS1155	BaFe	15.0 TB	360 MB/sec	32	7,680	9.78 Gb/in^2

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

Future Data Recording Projections

Commercially available tape areal densities have reached 9.78 Gb/in^2 while HDD areal densities are in 1,300 Gb/in^2 range, 10 times denser. Future density scenarios (see the most recent INSIC 2016 chart) project that annual HDD areal density growth rates will *not* maintain their historical 35-40% values and will likely slow toward 16%. Tape areal density growth rates are expected to double the HDD rate

averaging about 33% annually. The surface area available to increase HDD capacity on disk platters is crowded while tape cartridges have a surface area over 200 times greater than HDD to work with.

To increase capacity, many HDDs have increased the number of platters from three to seven while using helium filled disk enclosures to reduce friction. Modern tape drives operate at areal densities that are more than two orders of magnitude lower than the latest HDDs. It should therefore 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 are facing.

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 total capacity storage advantage for tape technology. NAND flash memory has surpassed hard disk drive (HDD) technology in areal density for the first time with announcements of 2,770 Gb/in^2. The transition to 3D-flash with 64-72 layers is imminent and SSD continues to apply pressure to the HDD industry. 3D XPoint will follow 3D Flash into the SSD market.

Tape Data Rates Soar

The INSIC 2015-2027 International Magnetic Tape Storage Roadmap projects tape data rates to increase at 22.5% annually and are expected to be as much as five times faster than HDDs by 2025. This is great news for businesses whose archives include data from the IoT and big data analytics, mobile and social systems, as well as content from higher density video streaming, hybrid cloud workloads and traditional data center applications. The faster data rates will significantly increase the benefits of RAIT, which allows multiple tape drives to transfer data in parallel providing a data rate multiplier.

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

Tape Roadmaps - LTO Roadmap Extended to Generation 12

The LTO Consortium publishes a well-defined roadmap (see below) and in Oct. 2017 extended the LTO family to LTO-12. With each successive LTO generation scheduled to arrive in approximately two and one-half year intervals, LTO continues to be the primary tape technology as it steadily improves the acquisition price, capacity and performance over previous models. In addition, 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. To push the innovation and capacity boundaries of LTO going forward, the current LTO format required a recording technology transition that supports high capacity growth for future LTO generations. As a result, the LTO-8 specification is only backwards compatible with the former generation LTO-7 rather than two prior versions. The new 12 TB native LTO-8 cartridges can hold 7,140,000 photos, 8,000 movies or 2,880,000 songs.

In addition, a new **LTO-7 Type M media** was introduced along with the launch of LTO-8 tape drives. LTO-8 drives can initialize new unused LTO-7 cartridges, yielding the LTO-7 Type M media, to an increased capacity of 9 TB native and 22.5 TB compressed instead of 6 TB and 15 TB in a standard LTO-7 drive. The

new LTO-7 Type M is manufactured and logoed as LTO Ultrium 7 media but labelled with a barcode label ending with the last 2 characters “M8”. The Type M data rate is the same as LTO-7 at 300 MB/sec.

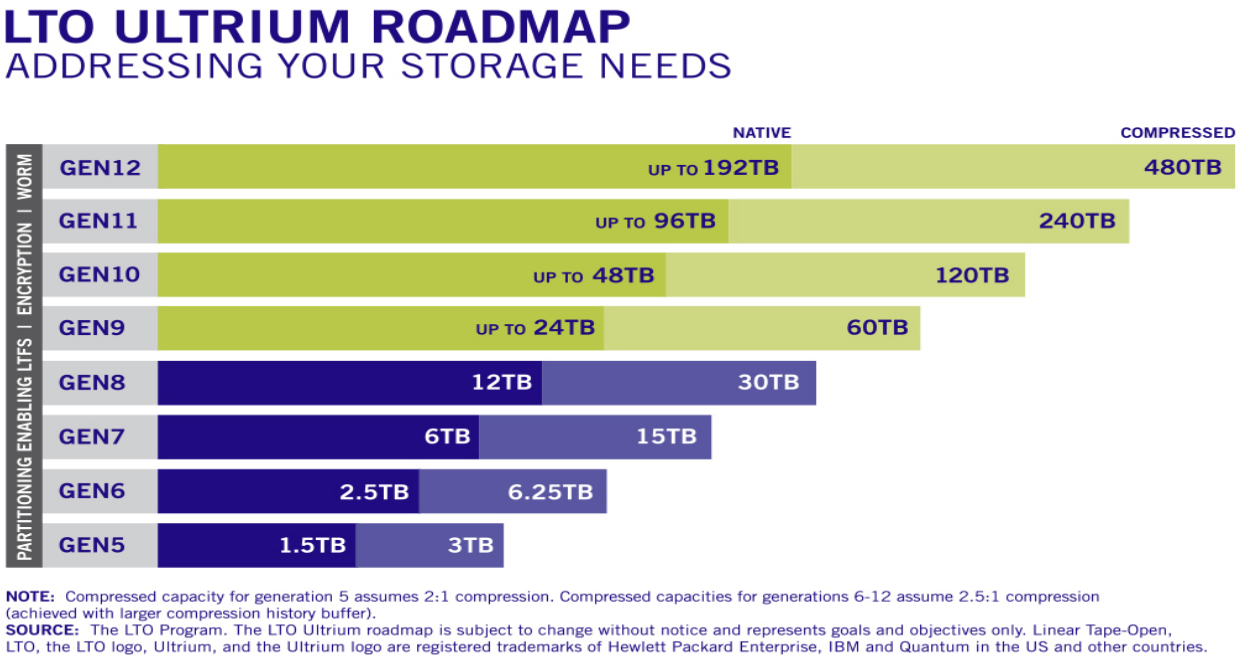

***Key Point:* The progress of future LTO tape systems is well defined, highly attainable, and is expected to support several more years of technology advancements. Expect similar improvements and progress for enterprise tape.**

Using Tape for Cybersecurity Prevention and the Air Gap

Security features are included on LTO and enterprise tape drives to address countless compliance and legal requirements including data encryption to protect data at rest, Write-Once-Read-Many (WORM), and various write-protect capabilities. Since tape is removable media, physical cartridges can be easily transported to another location in the case of extended power outages or natural disasters which have become more common in recent years. Given the number of causes that could potentially force data centers to go without electricity, media portability remains a final, but vital line of defense for data protection. As a result, the traditional “truck access method” has considerable value for data resiliency.

The tape air gap has ignited new interest in backing up data on tape. The “tape air gap” means that there is no electronic connection to the data on the removeable tape cartridge preventing a malware attack. Disk systems remaining online are always vulnerable to an attack. Backup data using the tape air gap is quickly regaining popularity as a key part of a comprehensive data security plan. The U.S. Department of Justice reports that an average of 4,000 ransomware attacks daily have been taking place since Jan. 1, 2016 and it’s now estimated that the global cost for affected businesses was nearly $5 billion by the end of 2017. In Sept. 2017, The Wall Street Journal published an article highlighting tape’s unique new role in cybersecurity.

Disk Challenges are Mounting - The Storage Squeeze Play Has HDD's Caught in the Middle

HDDs are increasing in capacity - but *not* in performance - as the IOPS (I/Os per Second) for HDDs have basically leveled off. The potential for more concurrently active data sets increases as HDD capacity grows and the increased contention for the single actuator arm causes further response time delays. Excessive RAID rebuild times are a major concern and it can now take several days to rebuild a failed high-capacity HDD impacting performance. As HDD capacities continue to increase, total time required for the RAID rebuilding process will become prohibitive for many IT organizations. Higher capacity HDDs could force a replacement for the traditional RAID architecture with such features as Erasure Coding. A notable shift in the storage landscape is underway as high-performance data moves from HDD onto flash SSD while lower activity and archive data migrate from HDD to modern tape. HDD shipment growth rates have declined nearly 35% since its highpoint in 2013 as the squeeze play is underway.

Archive Capability	Tape	Disk
TCO	Favors tape for archive as much as 6-15x over disk and cloud	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 years for most HDDs before upgrade or replacement, 7 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 cartridge has been demonstrated	Slowing capacity growth as roadmaps project disk capacity to lag tape for foreseeable future
Scale capacity	Tape scales by adding cartridges	Disk scales by adding more drives
Data access time	LTFS, the Active Archive and RAO improve tape access time	Disk is faster than tape for initial access and random-access apps
Data transfer rate	360 MB/sec for TS1155 and LTO-8, 360 MB/sec for LTO-8, RAIT multiplies tape data rates	Approx. 175 MB/sec for disk
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

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.***

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, with more energy and cooling. Tape holds a compelling value proposition advantage over disk for environmental requirements and data longevity. Some excellent TCO studies are publicly available and show the TCO for HDDs is approximately 6 to 15 times higher than for the equivalent capacity tape systems. See the easy to use TCO calculator from Brad Johns Consulting.

Emerging Applications Accelerate Future Tape Growth

Tape has been expanding its historical role as a backup solution to address a much broader set of requirements including data archives and disaster recovery services. Digital archives consisting of unstructured data, digital images, multi-media, video, social networks, compliance and surveillance data are the fastest growing data category experiencing a CAGR (Compounded Annual Growth Rate) in the range of 60% annually. Just ten years ago, large businesses generated roughly 90% of the world's digital data. Today an estimated 75-80% of all digital data is generated by individuals rather than by large businesses, however most of this data will eventually wind up in a large business or cloud provider's data center requiring unprecedented archival requirements. Much of this data will be destined for tape.

The Big Data Era is here, and the value of the digital archive is increasing as the benefits of mining and analyzing very large datasets enable analysts to project new business trends, prevent diseases, improve security, address cybercrime, and strengthen national defense, and much more. Presenting an ever-moving target, the boundaries of digital archives now reach petascale ($1x10^{15}$) or exascale ($1x10^{18}$) and will approach zettascale ($1x10^{21}$) capacity levels in the foreseeable future. Meeting these storage requirements only with disk and/or SSD will become financially prohibitive for most businesses.

Key point: ***The TCO of tape over disk is compelling for secure archival storage. With reliability having surpassed HDDs, tape is well positioned to capture the exploding demand for long-term storage.***

LTFS Enables Faster Data Access and Object Storage Support

Developed by IBM and introduced with LTO-5, LTFS (Linear Tape File System) provides an easier way to access and archive data to tape. LTFS has implemented 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 file-name searching using familiar drag and drop techniques are performed more quickly in server memory and do not require physical tape movement.

On July 5, 2017 LTFS announced connection with OpenStack Swift to enable movement of cold (archive) data from object storage to more economical tape storage for long-term retention. 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 data easier than ever before for tape applications.

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

RAIT Provides Much Higher Transfer Rates and Access Times Improve

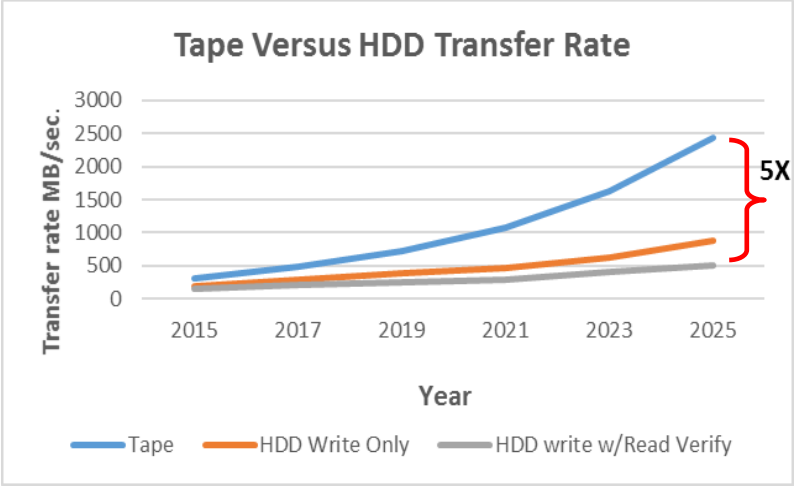

RAIT (Redundant Arrays of Inexpensive Tape) is available with HPSS (High Performance Storage System) and aggregates bandwidth across multiple tape drives in parallel significantly increasing data transfer rate (throughput). RAIT uses multiple tapes loaded in parallel for writing and reading data and provides parity for data reconstruction like RAID does for HDDs. Interest in RAIT is increasing as it takes advantage of significantly higher future tape transfer rates which are projected to yield tape data rates 5x faster than HDDs by 2025. The much higher transfer rates position RAIT for the HPC, hyperscale, cloud and enterprise markets. See INSIC data rate projection chart for tape and HDDs (above).

The tape industry has stepped up its focus on **access time** to a tape file by delivering 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 to reduces file access times as much as 50% while significantly reducing physical tape movement and wear. This capability becomes 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 movement times to optimally locate a tape cartridge which has the added benefit of improving library reliability.

Energy Consumption - Tape Means Green Storage

A commonly stated objective for many CIOs today is that *"if data isn't used, it shouldn't consume energy".* In response to this directive, the movement of archival data from HDD to more reliable, much more energy efficient, and more cost-effective tape storage is actively underway. Unlike storage providers, energy providers have shown little interest in lowering their rates (price per unit) and average data-center energy costs are growing at 10-20% per year or more per unit consumed.

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. Energy costs for tape capacity are typically less than 5% of the equivalent amount of disk capacity. The limits of power distribution in many data centers is being approached, 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 normally a last resort and is extremely expensive mandating that energy consumption be properly managed. Average IT electrical consumption rates for data centers are summarized in the chart below.

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 support 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.

Utility companies can restrict the amount of power some businesses can use at certain times of the day making their data center energy management strategy more critical. Hyperscale data centers such as Amazon, Apple, Facebook, Google and Microsoft face enormous energy management challenges which encourage using tape for storing enormous quantities of less active data.

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

The Active Archive Emerges to Improve Performance.

Tape's favorable economics for storing archival data are fueling increased interest in Active Archive solutions. An active archive is a combined solution providing users an automated means to store and manage all their archive data by integrating HDDs, SSDs, and tape in the data center or cloud. The active archive greatly improves tape access time by serving as an HDD or SSD cache buffer for a tape library. The active archive enables a high percentage of accesses to the tape subsystem to be satisfied from SSD or HDDs (cache hit ratio) improving access time to first byte of data.

***Key point:* The active archive provides disk-like cache buffer with faster access times for archival data.**

Tape Plays Increasing Role in the Cloud Era

Tape's role in the cloud provider market is expanding as tape is being recognized as the optimal solution for archival storage to address and reduce the higher HDD cost issues. Using HDDs for archival storage is a strategy – just not a very cost-effective one - and the cloud providers now understand this as their storage requirements soar. Since tape media is portable, using tape for cloud storage becomes highly advantageous if the cloud provider shuts down or should you want to quickly move your entire archive set to another cloud provider. Moving large amounts of archival data using available network bandwidth can take days or even weeks and can become cost prohibitive compared to moving tape via truck or airplane.

Key point: ***Cloud providers are deploying tape for their lowest-cost, most secure, long-term archival storage offerings. Storing archival data in the cloud represents a significant future growth opportunity for tape storage providers and a much lower cost solution for cloud providers.***

Summary

The magnetic tape industry has made significant progress in many areas over the past 10 years surpassing disk in many categories. Clearly the continued role for disk is established but disk is facing growing technological and cost challenges making capacity increases and any additional disk performance gains questionable. Tape will not replace HDDs or SSDs, but it has expanded its position as a highly cost-effective complement to SSD and HDD for the foreseeable future due to its lower TCO, high reliability, higher capacities, faster data rates, and significantly lower energy costs. Because of this progress, the tape industry is quickly re-positioning itself to address many new high capacity, long-term archival and big data storage repositories which now represent more than 60% of the world's total stored digital data. The rich technology improvements in the tape industry suggest that tape will continue to be the most cost-effective storage solution for the enormous archival and big data challenges that lie ahead, whether on-site, at a remote location, or in the cloud.

Bottom-line: ***The tape renaissance is positioning tape for the digital storage intensive world that lies ahead. For the tape industry - it's clearly a new game with new rules!***

End of report