

2021 Technology Update Series


TAPE.

NEW GAME. NEW RULES.

Tape Becomes a Key Enabler for the Zettabyte Era

Modern tape technology advancements are accelerating along with the digital transformation as we have now entered the era of zettabyte scale storage demand on an annual basis. The arrival of the zettabyte era has fueled enormous storage capacity demand pushing more and more businesses closer to exascale storage requirements. This amount of data quickly becomes cost-prohibitive to maintain on HDDs since 60-80% of all data stored is infrequently accessed and can be classified as archival or cold data. Tape has re-architected itself for the zettabyte era with the arrival of modern digital data tape technology. 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.

For tape it's clearly a new game with new rules!



FROM THE BEGINNING TAPE TECHNOLOGY PROGRESS HAS BEEN RELENTLESS

The tape industry has made tremendous strides in capacity and throughput since its inception. Consider the first successful magnetic tape drive arrived in 1952 having a media capacity of 2 megabytes per round reel and a data rate of 7,500 characters (bytes) per second. The latest tape drive has a media capacity of 20 terabytes (1x10¹²) and a data rate of 400 megabytes per second. Tape media capacity has increased 10 million times while the data rate has increased 53,333 times. The first successful disk drive arrived in 1956 with a capacity of 5 megabytes and a data rate of 10,000 characters per second. The latest disk drive has a capacity of 20 terabytes and a maximum data rate of 260 megabytes per

second. Disk drive capacity has increased 4 million times while the data rate has increased 26,000 times. The relentless technological progress of the tape industry has been driven by the need for highly secure data backup and the soaring amount of data that needs to be stored indefinitely. Keeping pace with storage demand is the perpetual challenge facing the entire storage industry. Fortunately, tape and disk have been joined in recent years by Flash SSDs (Solid State Disks) to address and stay ahead of the storage demands of the zettabyte era. Each technology has its unique considerations, and they combine to play a key role for optimizing data center storage.

LTO TAPE ROADMAP ADVANCES TO GENERATION 12

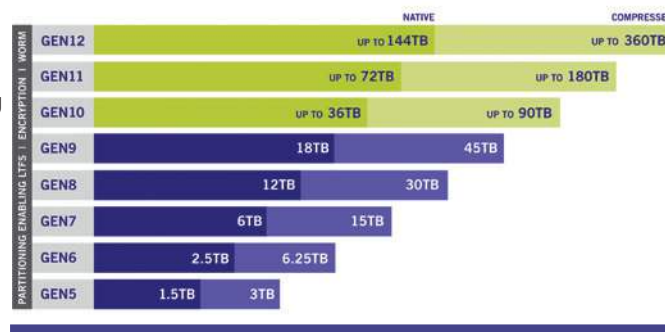
The LTO Consortium publishes a well-defined tape roadmap (see below) defining the LTO family through LTO-12. Each successive LTO generation is scheduled to arrive in approximately two- and one-half-year intervals, steadily improving the acquisition price/ TB, capacity, and performance over previous models. The newest LTO-9 tape technology will surpass the native capacity of an LTO-8 cartridge by 50 percent to 18 TB (45 TB compressed) and increasing drive throughput up to 400 MB/s enabling a single LTO-9 tape drive to write up to 1.44 TB/hour. A native 18 TB LTO-9 cartridge can hold 1,606,500 photos, 18,000 movies or 6,480,000 songs. The capacity increases set the stage for future tape growth as archival data continues to pile up much faster than it is being analyzed. The Covid-19 pandemic is a great example of this as it is generating medical data that will sit untouched for many years before it is fully analyzed. Given this trajectory, modern tape has become the leading strategic and lowest-cost storage solution for massive amounts of backup and archival data offering an additional highly secure, air-gapped layer of defense against cybercrime attack.

Beyond the capacity and throughput increases, the magnetic tape

industry has successfully re-architected itself delivering innovative technologies and functionality including vastly improved bit error rates yielding the highest reliability of any storage device, a media life of more than 30 years, and the lowest carbon footprint and TCO of any storage solution. 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 from IBM. The LTO-7 format expanded the “history buffer” in the drive compression engine, giving it a 2.5:1 compression ratio, up from 2:1 on previous LTO drives while the IBM enterprise TS1160

drive offers a 3:1 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 backwards compatible with the former LTO-7 generation rather than two prior versions. Additional tape functionality including LTFS, RAIT, RAO, TAOS, smart libraries and the Active Archive architecture adds further benefits to the core tape lineup. As the zettabyte era unfolds, these capabilities will have much greater significance.

LTO ULTRIUM ROADMAP
ADDRESSING YOUR STORAGE NEEDS



KEY POINT

The roadmap for future LTO tape systems is well defined, highly attainable, and is expected to support innovative advancements for the foreseeable future while protecting investments in tape automation infrastructure.

AREAL DENSITY IS THE 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 steady \$/TB 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) at right shows the tape and disk areal density roadmap progress and demos from 1989 and projects tape growth through to 2029. Lab demonstrations for tape technology are particularly promising for addressing the zettabyte era demands. Current LTO and enterprise tape drives operate at areal densities about two orders of magnitude lower than the latest HDDs. Therefore, the additional unused tape-recording area makes it possible to advance tape densities at historical rates for at least the next decade.

In April 2015, IBM and Fujifilm demonstrated a 220 TB (native) tape with 123 Gb/in² areal density using Barium Ferrite (BaFe) media. In December 2017, IBM and Sony demonstrated a 330 TB tape, using sputtered media with an areal density of 201 Gb/in². In Dec. 2020, IBM and Fujifilm demonstrated a record tape areal density of 317 Gb/in² yielding the potential for a 580 TB cartridge using a new magnetic particle called Strontium Ferrite (SrFe).

The LTO-9 18 TB tape has 12 Gb/in² areal density. An 18 TB disk drive has an areal density of 1,022 Gb/in². LTO-9 achieves the same 18 TB capacity with 1/85th of the areal density of an 18 TB disk. This provides significant tape development head room for tremendous growth. The 580 TB cartridge could store data equivalent to

120,000 DVDs which would stack taller than Burj Kalifa, the world's tallest building at 2,722 ft. Imagine 580 TB of data on a tape cartridge in the palm of your hand.

**317 Gb/in² demonstrates the sustainability of the INSIC Tape Roadmap
34% CAGR in Areal Density for the next decade**



The latest enterprise TS1160 tape drive using 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. Tape has a steeper growth rate, currently at 34 percent a year, than HDDs, which has a forecast areal density growth of 7.6 percent a year. With steadily increasing areal density capability demonstrated, expect tape to maintain its cost advantage vs. HDD and other technologies for the foreseeable future.

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 without any fundamental technology roadblocks.

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
LTO-9 BaFe	2020	18.0 TB (2.5:1)	400 MB/sec	32	8,960	12.0 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

BEYOND THE LTO ROADMAP PETABYTE TAPE CARTRIDGES ARE ON THE HORIZON

Fujifilm is developing future magnetic particles that could lead to the achievement of one petabyte of capacity on a single tape cartridge that far exceeds the highest capacity on the LTO roadmap, which currently extends out to LTO-12 and 144 TB. Today’s Fujifilm tape uses a Barium Ferrite (BaFe) magnetic layer. The company is expected to follow this with Strontium Ferrite (SrFe) media, which have smaller nanoparticles that are less than 60 per cent of BaFe particle size. IBM and Fujifilm have already demonstrated a 580 TB raw capacity tape using Strontium Ferrite media. For the 1 PB tape cartridge development, Fujifilm is developing advanced technology based on Epsilon Ferrite ($\epsilon\text{-Fe}_2\text{O}_3$), which has even smaller nanoparticles. It should be possible to continue scaling tape

areal density at historical rates beyond the next decade providing a sustained volumetric and native capacity advantage for tape technology.



TAPE LEADS STORAGE RELIABILITY RATINGS

Reliability is the critical lifeline for digital storage. For tape, customers have indicated for years that a key cause of failure was due to media and handling errors; however, these concerns have since been addressed. HDD reliability concerns are mitigated with high availability features based on RAID and various types of replication. Tape reliability concerns are mitigated by creating multiple copies. For all practical purposes, “two copies on tape” could be considered the equivalent of mirroring data (RAID 1) on HDDs. Specifically for tape, 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). When tape is being read, it is streamed over the head at

a speed of about 15km/h and with the new servo technologies, the tape head can be positioned with an accuracy that is about 1.5x the width of a DNA molecule.

To improve reliability, LTO drives switched to PRML from the older RLL (Run Length Limited) error checking code. 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 enabling a much higher recording density while enabling 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. Today, both LTO and enterprise tape products are more reliable than any HDD. Times have changed!

STORAGE DEVICE RELIABILITY RATINGS Source: Vendor’s published BER	BER (BIT ERROR RATE) Bits read before permanent error
Enterprise Tape TS1160	1×10^{20} bits
Enterprise/LTO Tape (TS1155, LTO-7, 8, 9)	1×10^{19} bits
Enterprise SSD (NAND)	1×10^{18} bits
LTO 5-6	1×10^{17} bits
Enterprise HDD (FC/SAS)	1×10^{16} bits
Enterprise HDD (SATA)	1×10^{15} bits
Desktop/consumer HDD (SATA)	1×10^{14} bits

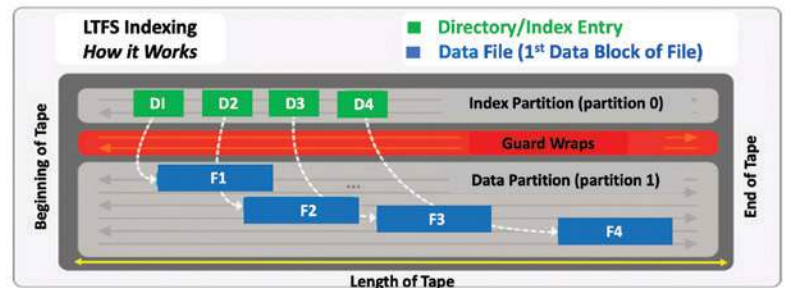
KEY POINT

Tape has the highest reliability of any data storage device and even higher BER levels are projected in the future.

LTFS ENABLES FASTER DATA ACCESS FOR FILES, OBJECTS, AND ARCHIVES

Developed by IBM and introduced in 2010 with LTO-5, LTFS (Linear Tape File System) provides a significant step forward in moving tape storage away from its reputation as complex and difficult to use providing 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. LTFS also makes retrieving object storage, commonly used for archival data, much easier for

tape applications to use. LTFS continues to gain momentum as 33 companies are LTFS implementers. Thanks in large part to LTFS, LTO has become the de-facto standard in the Media & Entertainment industry for file sharing and archival applications.



KEY POINT

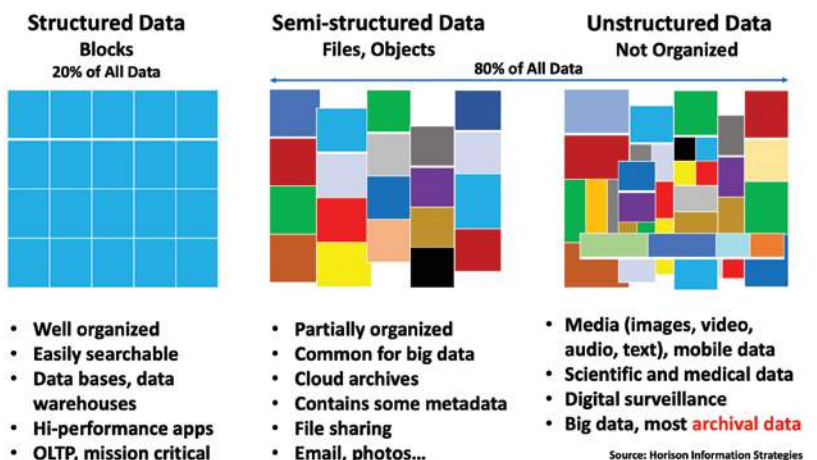
LTFS partitioning, and its future iterations will provide even greater access capabilities for tape and attract more ISVs (Independent Software Vendors) to exploit its capabilities.

OBJECT STORAGE AND TAPE COMBINE TO ADDRESS SOARING ARCHIVAL DEMAND

Understanding data formats has become increasingly important for effective data classification. Block storage chunks data into arbitrarily organized, evenly sized volumes typically stored in databases; file storage organizes and represents data as a hierarchy of files in folders; and object storage manages data and links it to associated metadata. Object storage evolved out of the need to optimize storage performance, improve scaling capabilities, and to facilitate data analytics through its use of metadata for very large volumes of unstructured (not organized) data. Sem-structured and unstructured data accounts for about 80% of all stored data.

requirement, Fujifilm has developed the Fujifilm Object Archive software that enables object storage on tape systems giving users the opportunity to store shared data in the cloud (and hybrid cloud) while storing large-capacity cold data on low-cost, long-term tape storage. Once this content is archived, Object Archive software provides a secure, long-term archival copy on tape, with a physical air gap providing immutable data copies that can help prevent a ransomware attack.

Growth of semi-structured and unstructured data has become one of the biggest data center storage challenges as object storage requirements can quickly scale to hundreds of petabytes in a single namespace. Object storage has surpassed the growth rate of conventional file and block storage formats, growing over 30% annually becoming the preferred cloud storage format and that trend looks to continue. Software-defined object storage solutions using modern LTO and 3592 enterprise tape have emerged to provide a more economical solution than using traditional file and block storage systems based on HDD or flash technology. To address this growing storage



THE TAPE AIR GAP PROVIDES SECURITY AND CYBERCRIME PREVENTION

Air gap data storage, 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 removable tape cartridge therefore preventing a malware attack on stored data. HDD and SDD systems remaining online 7x24x365 are always vulnerable to a cybercrime attack. Across the world, hackers are exploiting numerous security weaknesses and holding the data of organizations hostage, sometimes demanding tens of millions of dollars in payment. Ransomware attacks on the Colonial Pipeline (\$4.4 M), JBS Foods (\$11.0 M), and other major organizations have made headlines in 2021, and show no sign of slowing down. Since the barriers of entry are extremely low, expect to see more attacks across every industry as damages are projected to soar. Cybersecurity Ventures predicts ransomware will cost the world approximately \$265 billion annually by 2031 with a new attack occurring every 2 seconds as ransomware attackers refine their malware payloads and related extortion activities. The dollar figure is based on 30 percent year-over-year growth in damage

GLOBAL RANSOMWARE DAMAGE COST

2015: \$325 MILLION
2017: \$5 BILLION
2021: \$20 BILLION
2024: \$42 BILLION
2026: \$71.5 BILLION
2028: \$157 BILLION
2031: \$265 BILLION



Ransomware is expected to attach a business, consumer, or device every 2 seconds by 2031, up from every 11 seconds in 2021.

Source: CYBERSECURITY VENTURES

costs over the next 10 years. Air gapping your data should now be an integral part of any archive, backup, and recovery strategy whether on-premises or in the cloud and has made tape a key strategic component of modern backup strategies. It's no longer a matter of “if” but “when” hackers will breach your network.



TAPE PERFORMANCE IMPROVEMENTS BOOST ACCESS TIMES AND THROUGHPUT

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, CSP and enterprise markets. RAIL (Redundant Arrays of Independent Libraries) is poised to arrive as a component of geo-spreading and similar to RAIT, stripes data across different tape cartridges but in different libraries. RAIL can provide increased geographic resiliency and availability in case of full data center outage or natural disaster and can create higher availability archives for geo-spreading.

The tape industry has stepped up its focus to improve tape file access time with two features, 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 can improve 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.

KEY POINT



Tape technology has delivered features that provide faster access to the first byte of data, much faster aggregate data rates, and intelligent libraries that minimize robotic arm movement.



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 enterprise and hyperscale data centers strive to keep servers cool and operational. Data centers and information technology currently consume roughly 3% of the world’s electricity 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”. Tape’s low cost per TB, Total Cost of Ownership (TCO) and carbon footprint advantages compared with other storage mediums has established it as the most cost-effective storage technology for long-term, secure data retention. CO₂ and TCO reduction favors tape by as much as 87% and 86% respectively over HDDs (see charts at right). 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. The massive energy footprint of data centers has become a force for climate action over the past decade, with tape and cloud computing delivering the greatest advances in sustainable operations. According to a new study *Improving Information Technology Sustainability with Modern Tape Storage* by Brad Johns, tape produces 95% less carbon emissions compared to HDDs over its complete lifecycle from raw materials to disposal.

The business world has responded to the sustainability challenge. According to the U.S. Chamber of Commerce, over 90% of the S&P 500 companies now publish sustainability reports. The limits of power consumption in many data centers, especially hyperscale, are being reached forcing



Source: Brad Johns Consulting, for 10PB, 35% CAGR, 10 years

organizations to explore new cooling techniques such as water-cooled racks, outdoor and mobile cooling, or in some cases, building an additional data center. Building another data center is a last resort and is extremely expensive mandating that energy consumption always be efficiently managed. Best practices for reducing data center energy consumption 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 representing a real opportunity for data centers to significantly reduce their carbon emissions and overall impact on the environment. The benefit of moving cold data from HDDs to modern tape storage cannot be underestimated.

KEY POINT

Total cost of ownership and CO₂ emissions heavily favor tape over disk. Shifting less active and archival data from disk to tape storage and virtualizing servers are the most significant ways of reducing energy consumption expenses in the data center.



COMPARING TAPE AND DISK FUNCTIONALITY

FUNCTIONALITY	TAPE	DISK
TCO and CO₂ Reduction	Favors tape by as much as 86% and 87% respectively over HDDs	Much higher TCO and carbon footprint, more frequent conversions and upgrades
Long-life media	More than 30 years 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) has reached 1×10^{20} versus 1×10^{16} for HDDs	Disk BER not improving as fast as tape
Inactive data does not consume energy	True for tape, this is becoming a goal for most data centers. "If the data isn't being used, it shouldn't consume energy"	Rarely true 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 attacks	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 580 TB tape capacities 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 file access time	HDD is much faster (ms) than tape (secs) for initial and random access
Data transfer rate	400 MB/sec for TS1160, 400 MB/sec for LTO-9, RAIT multiplies data rates	Approx. 160-260 MB/sec for HDDs
Portability - Move media for DR with or without electricity	Yes, tape media is completely removable and easily transported in absence of data center electricity	HDDs 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 & hyperscale 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

Even with the global Covid-19 pandemic still active and how it ultimately plays out remaining unclear, the digital data storage landscape continues to expand presenting a constant stream of new challenges while demanding continual innovation. 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 will capitalize on 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 large-scale 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. For tape it's clearly a new game with new rules.

Horison Information Strategies is a data storage industry analyst and consulting firm specializing in executive briefings, market strategy development, whitepapers and research reports encompassing current and future storage technologies. Horison identifies disruptive and emerging data storage trends and growth opportunities for end-users, storage industry providers, and startup ventures.

© Horison Information Strategies, Boulder, CO. All rights reserved.