HORISON
Information Strategies
Fred Moore, President
Horison.com

Sponsored By
FUJIFILM
Value from Innovation

STRATEGIC SOLUTIONS FOR **DATA STORAGE AT SCALE**

Leveraging Hybrid Storage Solutions to Solve Massive Data Storage Challenges

The onset of the zettabyte era, coupled with a growing wave of traditional enterprise and large scale data centers (LSDCs), had just preceded the arrival of the COVID-19 pandemic which quickly impacted virtually every aspect of human life as well as the IT industry. Connectivity to digital content became essential for many businesses to survive. As a result, select workloads shifted to the cloud—public, private, and hybrid—putting even more pressure on LSDCs to build the most efficient, highly scalable, and secure storage infrastructures possible. LSDCs can have thousands of servers and petabytes of storage to manage. ***Traditional storage management techniques often left both LSDCs and cloud service providers (CSPs) struggling with as much as 60% or more of their data stored on the wrong tier of storage wasting billions of dollars a year.*** The challenges that this shift presents becomes obvious.

LARGE SCALE DATA CENTERS FOCUS ON OPTIMIZATION

Modern enterprise and large scale data centers (LSDC) are different than they were just a short time ago. Key data center design components include storage, servers, routers, firewalls, and HVAC systems. The effective use of power, space, cooling and orchestrating dynamic, automated software management to minimize or eliminate downtime are crucial tenants for achieving ultra-efficient data center operations. Achieving maximum cost savings remains the top pain point for LSDCs and places a premium on storage TCO.

HYBRID CLOUDS ARE BECOMING INTEGRAL TO LSDC OPERATIONS

Many LSDCs are built, owned, and operated by companies for their end users and are often housed on the corporate campus. Whenever applications and data are hosted in the cloud, they are using data center resources from the CSP. Many LSDCs are combining their traditional on-premises physical servers and storage with public cloud systems to support select workloads in a hybrid cloud environment. The hybrid cloud has become the most popular form of cloud deployment for larger data consumers and customers as using the cloud is no longer an either/or choice. A hybrid cloud is a computing ecosystem that combines an on-premises data center (also called a private cloud) with a public cloud, allowing data and applications to be shared between them. A hybrid cloud can include "multicloud" configurations where an organization uses more than one public cloud provider in addition to their on-premises data center.

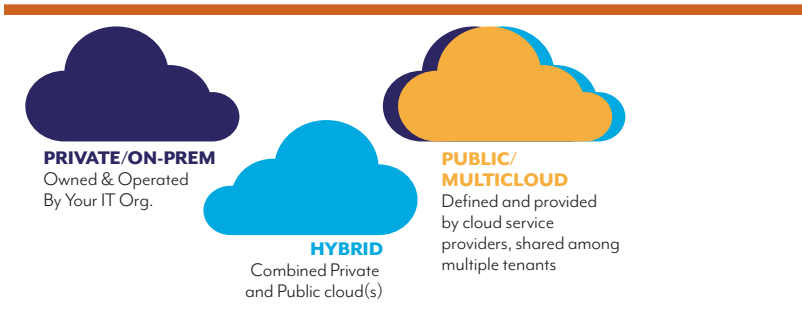

The Cloud Computing Trends: 2021 State of the Cloud report indicates that 92% of enterprises have a multi-cloud strategy and 82% have implemented a hybrid cloud strategy. This strategy includes a combination of storage technologies from flash, to HDD to tape that complement each other in a tiered storage environment.

STORAGE TECHNOLOGY PROGRESS IS RELENTLESS

The relentless technological progress of the storage industry has been driven by the soaring amount of data generated that needs to be stored. Keeping pace with storage demand is the perpetual challenge facing the storage industry. Fortunately, tape and disk have been joined by Flash SSDs (Solid State Drives) 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.

Fortunately the storage industry recognizes that cloud depends heavily on data storage solutuons and continues to make tremendous strides in capacity and throughput. Consider the first successful 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. Current LTO tape drives have a media capacity of 12.0 TB (LTO-8), and 18.0 TB (LTO-9) while IBM TS1160 drives can store 20 terabytes with 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.

Flash SSD Considerations – NVMe and Flash solid-state memories play a major role in providing high-performance for LSDCs, use less power and are easier to manage than HDDs. Though the price of flash-based SSDs is closing the gap with HDDs, cost is typically not the major selection factor for SSDs as they remain the most expensive storage device on a $/TB basis. Flash SSDs have a Flash Translation Layer (FTL) intelligently managing everything from caching to performance, wear leveling, to garbage collection, etc. Any flash overhead will be tolerated as the overall performance gains compared to HDDs are compelling.

HDD Considerations – HDDs are the mainstay of the storage environment but are facing challenges. Traditional RAID data replication architectures can become too expensive, too slow, and unmanageable. The higher the HDD capacity, the longer the RAID rebuild time required to restore or recover a failed drive to regain redundancy. A 4 TB HDD will take at least 10 hours to rebuild – a 20 TB HDD can take several days or even weeks which are unacceptable timeframes for most data centers. With RAID reaching its practical limits, except for higher IOPs data, Erasure Coding has emerged as an availability alternative to RAID in which data is broken into fragments (shards) that are stored across different geographical locations (geo-spreading). After servers, HDDs are the largest consumers of data center energy.

Tape Considerations – Typically, 60-80% or more of all data is either never or infrequently accessed after 90 days and is classified as archival making it the largest data category. As a result, automated tape library usage in LSDCs is on the rise to easily scale capacity and to contain out of control HDD growth by storing colder data and unstructured data. Tape provides the lowest storage TCO, carbon footprint and energy costs and provides hacker-proof cybercrime security via the tape air gap making it the ideal archive technology. As a result, several CSPs now offer specific archival and cold storage services based exclusively on removable tape cartridges.

Storage Refresh and Upgrade Cycle Considerations – The frequency of hardware upgrades can become labor-intensive and presents a key consideration for operations. The refresh cycle must be a seamless, non-disruptive process. The useful life for HDDs typically lasts an average of 4-5 years before replacement. Tape drives typically last 8-10 years, tape libraries last 10-20 years and modern tape media life is rated at 30 years or more. Fortunately, over 60% of the world's stored data is most cost-effectively stored on tape where refresh cycles are least frequent.

THE VALUE OF TAPE INCREASES AS DATA CENTERS GROW

Today LSDCs are leveraging the many advantages of tape technology solutions to manage extreme data growth and long-term retention challenges. Archival data is the fastest growing storage classification and requires secure, long-term storage solutions for regulatory reasons or due to the potential value that the data can provide through interpretation or analysis. Modern tape architectures enable LSDCs to achieve their data preservation objectives by providing backup, recovery, archive, easy capacity scaling, the lowest TCO and carbon footprint, highest reliability, the fastest throughput, and cybersecurity protection via the air gap. The value of these benefits is expected to increase for tape technology going forward.

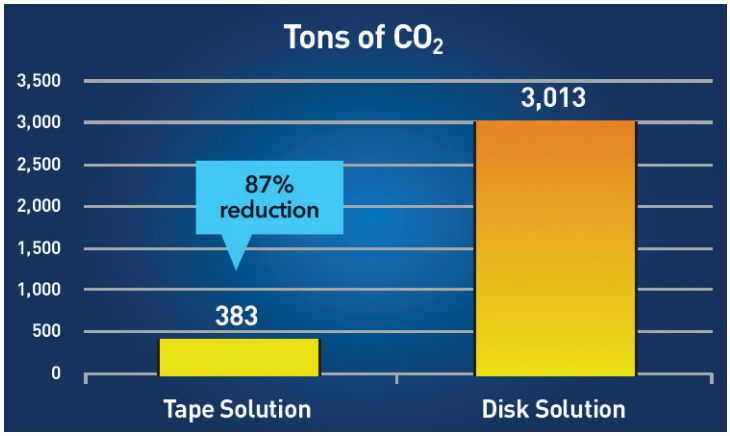

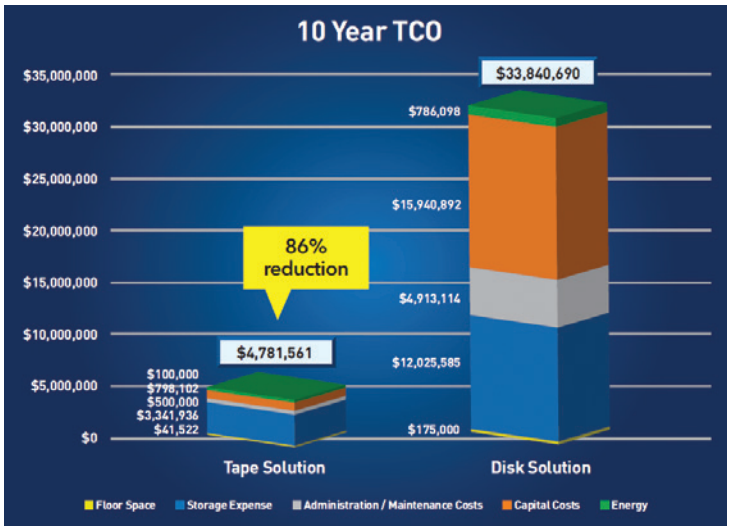

Source: Brad Johns Consulting, for 10PB, 35% CAGR, 10 years

How a CSP Leverages Tape

CASE STUDY

This CSP has more than 20 major data centers across the globe with exabytes of online data. LTO based digital tape technologies are deployed to ensure all their data is backed-up, secure and recoverable. The challenges associated with protecting this sheer amount of data extends well beyond the technical, as cost becomes a major variable.

When managing at hyperscale, the best strategies are to employ the most efficient processes in both software and hardware including all the overhead costs associated with infrastructure and management. First and foremost, this CSP started with a few company-wide principles that guide all projects and their data sets with the first being that all data must be backed-up. Given this principle and the hyper-scale of their data, this becomes a seemingly daunting task. However, if the data is not worth backing up, the project is not worth doing!

The second principle applied by this CSP is that all back-ups must be tested by restoring samples of the data. This principle ensures that issues with recovery of data are found before the data needs to be recovered. While these guiding principles carry a considerable cost, this burden is far outweighed by the damage of unrecoverable client data.

Data protection systems at this CSP are distributed across multiple regions for fault tolerance and geographic protection with large enterprise LTO tape libraries on each site. Data is sharded utilizing 20+8 Reed Solomon erasure coding with 40% protection overhead. Data is written to tape using RAIT 4+1 (similar to RAID 4) with data written to 4 tape drives at once. This CSP utilizes an archive process to retain data for longer periods of time with archive retention periods set by the client. The deployed tape technology provides the longest interval between required migrations further reducing TCO by reducing migration processes and forklift upgrades.

The efficiency of this CSP's tape-based data protection and archive strategy has resulted in a more cost-effective and highly reliable solution. This solution provides greater recoverability while protecting against multiple failures with a total of 75% overhead compared to 100% for a single copy of data or 200% for two copies.

LABORATORY DEMONSTRATIONS SHOW FEW CAPACITY LIMITS

In December 2017, IBM and Sony demonstrated the potential for a 330 TB tape cartridge, using sputtered media with an areal density of 201 GB/in^2. In Dec. 2020, IBM and Fujifilm demonstrated a record tape areal density of 317 GB/in^2 yielding a future potential 580 TB cartridge using a new magnetic particle called Strontium Ferrite (SrFe). 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 removeable storage media. Tape has a steeper areal density growth rate, currently at 34% a year, compared to a disk drive which has a forecast growth of just 7.6% 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.

CARBON FOOTPRINT, ENERGY AND TCO ARE LOWER WITH TAPE

Data centers and information technology currently consume roughly 3% of the world's electricity and is expected to soar up to 8% by 2030. Data centers often use an amount of electricity that could power an entire city. The substantial electricity use by data centers increases the concerns about carbon dioxide (CO_2) emissions. A key factor in the reduced rate of CO_2 emissions has been the aggregation of smaller data centers to larger-scale facilities that can more efficiently manage power capacity, optimize cooling, and increase server utilization rates.

Nonetheless the energy demands in many data centers force them to continuously explore new cooling techniques. Studies indicate tape can provide a huge advantage providing TCO and carbon footprint reductions (by 86% and 87% respectively - see charts below) and promote sustainability. Note: The average US cost of electricity for Feb. 2021 was 13.9 cents per KwH. A commonly stated objective for many data center managers today is that *if data isn't used, it shouldn't consume energy*.

REMOVABLE MEDIA MOVES LARGE AMOUNTS OF DATA QUICKLY

For hybrid cloud LSDCs, reverse migration of data in or out of the cloud to and from on-prem storage using network bandwidth can take days to several weeks and can become cost prohibitive compared to physically moving tape media. Since tape media is removable (portable), tape becomes advantageous providing leverage if a cloud service provider shuts down or should you want or need to "quickly" move your entire petascale media archive to another provider in the event of a natural disaster or extended power outage. For example, it takes 31 hours and 6 minutes to transfer 1 PB of data at a 100 Gig E data rate. One PB requires just 83 LTO-8, 12 TB cartridges which can easily be loaded in a car, truck or airplane and be moved to a new location in a few hours or less depending on the distance travelled.

THE TAPE AIR GAP PROVIDES AN EXTRA LEVEL OF SECURITY FOR CYBERCRIME PREVENTION

Fighting the cybercrime epidemic has become a major focus for most data centers. Practically every form of data storage has the potential to be corrupted. Air gapped data storage, inherent with tape technology, has generated a 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 data stored offline on tape. HDD and SSD systems remaining online 7x24x365 are always vulnerable to a cybercrime attack.

In 2020, the average data breach cost was $3.86 million while the highest ever single ransomware demand grew to $30 million. High-profile ransomware incidents are on the rise and ransomware is projected to cost organizations globally as much as $20 billion by 2021. Air gapping data should be an integral part of any archive, backup, and recovery strategy whether on-premises or in the cloud.

The added built-in tape features of encryption and WORM provide even more security.

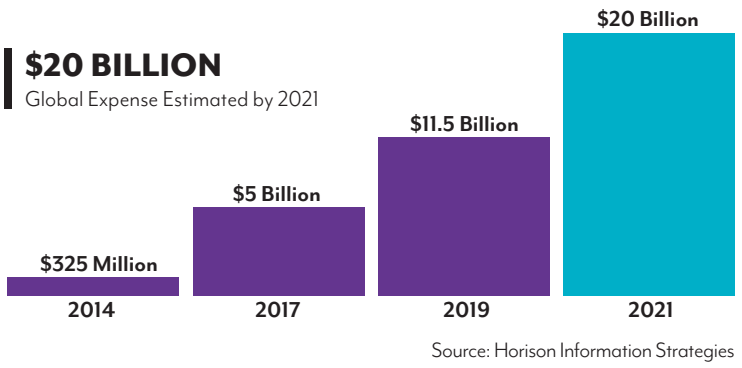

Source: Horison Information Strategies

TIERED STORAGE BECOMES ESSENTIAL FOR MAXIMUM STORAGE EFFICIENCY

LSDCs are taking advantage of tiered storage by more closely integrating high-performance SSDs, HDD arrays with automated tape libraries to optimize on-premise storage systems. Even though LSDCs and enterprises are struggling with the exploding growth of disk farms which can devour IT budgets and overcrowd data centers, many continue to maintain expensive disks, typically half full of data which can often have little or no activity for several years. Obviously, few data centers can afford to sustain this degree of inefficiency. The greatest benefits of tiered storage are achieved when tape is used. Tape's scalability, lower price, lower TCO and reduced carbon footprint provides greater benefits as the size of the storage environment increases. The magnitude of zettabyte era will sooner or later make tape mandatory for sheer economic survival. For the LSDC *adding disk is tactical – adding tape is now strategic*.

THE TIERED STORAGE MODEL

Data Classification and Optimal Allocation by Tier

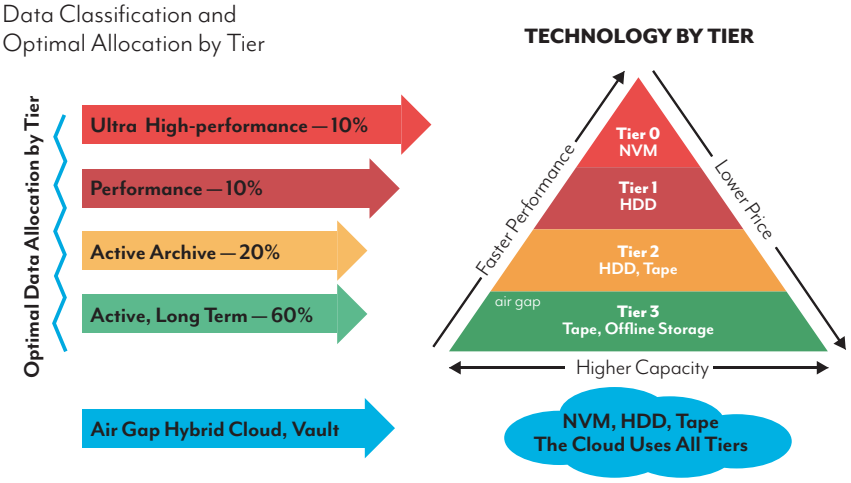

Source: Horison Information Strategies

TAPE SCALES BULK STORAGE CAPACITY EASILY

LSDCs require fast, easy scaling for storage capacity. One petabyte using 20 TB HDDs requires 50 drives and one exabyte requires 50,000 20 TB HDDs. Keep in mind scaling HDD capacity gets expensive quickly as each additional drive requires additional energy to spin and cool. Tape easily scales capacity by simply adding media without the requirement for adding more energy demand. For LSDCs *compute and storage must scale non-disruptively and independently of one another*.

TAPE PROVIDES THE MOST COST-EFFECTIVE SOLUTION FOR OBJECT STORAGE

Object storage has surpassed the growth rate of conventional file and block storage formats, growing over 30% annually becoming the preferred cloud archive storage format. Object storage presents a growing LSDC infrastructure challenge as requirements can quickly scale to petabytes in a single namespace. Structured data accounts for about 20% of all stored data, is well organized and is typically stored in databases in block format. Semi-structured and unstructured data accounts for about 80% of all data, is typically archival, not well organized and is stored in either file or object format and is well suited for tape. Software-defined object storage for LTO and 3592 enterprise tape has emerged to provide a more economical solution serving as a relief valve for low-activity object data sitting idle on HDDs. Over the next five years, large enterprises are projected to triple their amount of unstructured data stored on-premises, at the edge and in the public cloud.

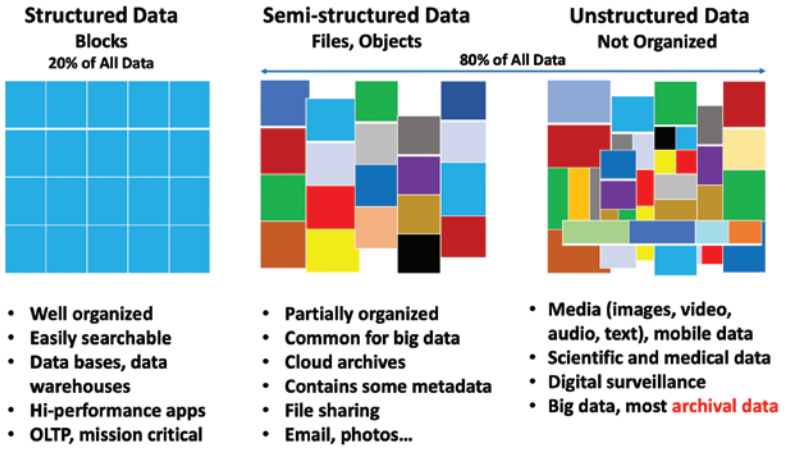

Source: Horison Information Strategies

INTELLIGENT STORAGE SOFTWARE

LSDCs rely on software-defined storage management solutions and are beginning to benefit from AI and ML delivering a higher degree of automation and self-healing while minimizing direct human involvement. LSDCs virtualize servers to control server sprawl and reduce energy consumption with as many operating system images as possible running on each physical server. AI will also improve automated data migration between storage tiers to move low activity data from HDDs to more cost-effective tape. Intelligent software and features including RAID, RAIT, replication and erasure coding are readily available to build higher-availability and fault tolerant storage architectures. In addition, the OCP (Open Compute Project) has an Archival Storage sub-project focusing on unique open-source archival and object solutions including tape that are more efficient, flexible, and scalable.

Function	Benefits Summary
Price/TCO/CO2	Tape has the lowest acquisition price $/TB, lowest TCO and lowest carbon footprint.
Performance	Much improved-active archives, fastest data rates, RAIT, smarter and faster robotics, time to 1st byte features (RAO, TAOS) have arrived improving tape access times.
Capacity	LTO-9 cartridge capacity @18 TB (45 TB compressed) with 400 MB/sec data rate. exabyte capacity libaries are available. Lab demos demonstrate tape capacities can reach 580 TBs.
Open Standards	LTO tape drives and LTFS provide open standard interface and acces.
Scalability	Tape scales capacity by adding media ***without adding*** energy consumption, HDDs add capacity by adding drives ***with adding*** energy consumption.
Energy/CO2	Tape uses much less energy and has much lower carbon footprint than HDDs (~85%)
Portability	Tape media easily portable, HDDs difficult to move
Cybersecurity	Tape air gap prevents cybercrime attacks, strong defense against malware.
Durability/Media Life	LTO reliability ($1x10^{19}$), has surpassed HDDs ($1x10^{16}$), media life>30 years for all modern tape.
Recording Limits	HDDs facing areal density and performance limits. Well defined and sustainable roadmap for tape.

DENSE RACKS OPTIMIZE FOOTPRINT

Rack management is a key discipline for many LSDCs and CSPs. Rack density is increasing as LSDC's squeeze servers, SSDs, HDDs and often power supplies directly into the larger racks as opposed to using stand-alone SANs or DAS, to achieve the smallest possible footprint. Tape libraries don't generate much heat, and some library vendors are taking packaging efficiency a step further by offering optional 5U or 10U top racks to reduce the overall footprint and simplify cabling. The extra rack space above the tape library can be used for power distribution units, I/O interface switches, data movers, and other hyper-converged nodes. Space can become a premium for LSDCs and eventually every cubic inch counts. To maximize space utilization, data centers will build racks as high as possible. Currently the 48U rack is the preferred choice and a 70U (122.5") open frame rack that pushes the boundaries of existing rack mounting height is available. Note: one rack unit (U) is 1.75".

COOLING TECHNOLOGY IS CRITICAL FOR ENERGY CONTAINMENT

For many LSDCs the largest operational expense can be their cooling systems. Higher power consumption and increasing carbon emissions have forced data centers to seek new energy sources to reduce and more effectively manage energy consumption. For years data centers were normally built with raised flooring. This was done to allow for airflow, power, chilled water, cable routing and provide more flexibility during moves and changes. Over time the underfloor area could become quite congested, and cool air delivery became erratic. For most modern LSDCs, concrete floors, or slabs, have become the more common choice, with cables and cooling systems running above the computer equipment rather than below offering a better job of routing cold air. For storage systems, tape libraries consume very little power and use a rack-row design that easily fits data center layouts and maximizes air handling and heat dissipation.

LOOKING AHEAD

Major Shifts in the Large Scale Data Center Landscape

- The growth of LSDCs is fueled by the migration of many smaller data centers to fewer, but larger and more efficient data centers.
- Total cost savings remains the top pain point for LSDCs placing a premium on storage TCO.
- Energy savings, carbon footprint reduction and sustainability are critical LSDC issues.
- Moving low activity data from HDDs to tape greatly reduces TCO and CO_2 emissions.
- Containing storage environments based solely on HDDs will become cost prohibitive.
- Tape scales by adding more media and HDDs scale by adding more drives which require more energy to spin and keep cool.
- The greatest benefits of tiered storage are achieved when tape is used.
- Tape's ease of removability and portability provides air gap protection against cyber criminals.
- As the zettabyte era accelerates, tape will become mandatory for data center economic survival.

CONCLUSION

The zettabyte era has pushed the limits of large-scale data centers while disrupting the traditional data center model. Some of the biggest challenges facing the LSDCs include TCO, extreme data growth, data security, carbon footprint, reducing energy demand, and they have all taken center stage. Today's LSDCs are re-engineering their on-premise and public cloud storage strategies to more cost-effectively manage data growth while taking advantage of the compelling economics and benefits of modern tape at scale. Without a new disruptive technology on the horizon to cost-effectively contain the world's archival and cold data, the many rich tape technology improvements of the past 10 years and the promising technology roadmap should continue to make tape the most cost-effective storage solution for the unprecedented data demands in the foreseeable future.

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.