Strategic Solutions for Hyperscale Data Storage

Leveraging Hybrid Storage Solutions to Solve Massive Data Storage Challenges

Horison Information Strategies
Fred Moore, President
www.horison.com

Introduction

The wave of HSDCs (Hyperscale Data Centers) had just preceded the onset of the zettabyte era when the COVID-19 pandemic arrived impacting virtually every aspect of human life as well as the IT industry. After the initial standstill, new rules emerged from the pandemic creating new norms for business, and new strategies for personal engagement. As the world sheltered in place often working remotely, some industries had to reduce their overall IT demand. The value of industries such as airlines and transportation, entertainment, hospitality, tourism, sports, retail, in-class education, attendee conferences, and logistics, is mainly created through in-person experiences. Many of these businesses could not easily transform all their offerings into digital content and their overall IT demand was affected.

Other organizations, where value is created through digital experiences, became almost 100% virtual transforming themselves by offering increased digital content delivery and services. Examples of these technology-based industries include e-commerce, e-medicine, digital media, telecommunications, video conferencing, e-learning, and a growing work at home digital workforce emerged. In these cases, the Covid-19 pandemic created surges in global digital demand and quickly pushed IT functions to HSDCs, CSPs (Cloud Service Providers), and large enterprises to immediately address unexpected growth. This has put even more pressure on HSDCs, CSPs and large enterprises on their way to hyperscale status to build the most efficient, highly scalable and secure storage infrastructures possible.

A Closer Look at Hyperscale Influence on the IT Landscape

Hyperscale refers to a computer architecture that scales massive compute power, memory, a high-speed networking infrastructure, and enormous storage resources typically serving millions of users with relatively few applications. While most enterprise data centers can rely on out-of-the-box infrastructures from vendors, hyperscale companies must personalize nearly every aspect of their environment. A HSDC architecture is typically made up of tens to hundreds of thousands of small, inexpensive servers, commodity components and nodes, providing enormous compute, storage, and networking capabilities. HSDCs use Artificial Intelligence (AI), and Machine Learning (ML) to automatically manage workloads and are leveraging the entire storage hierarchy including significantly increased tape usage to improve cost efficiencies and improve sustainability. Large-scale data centers (LSDCs) are becoming the next wave of hyperscalers. There were an estimated 570 HSDCs WW at year end 2020 accounting for 47% of all data center servers and 57% of all data center storage! The enormous impact of hyperscale computing in the global IT landscape going forward promises a transformative model for existing data center operations.

The Hyperscale Influence

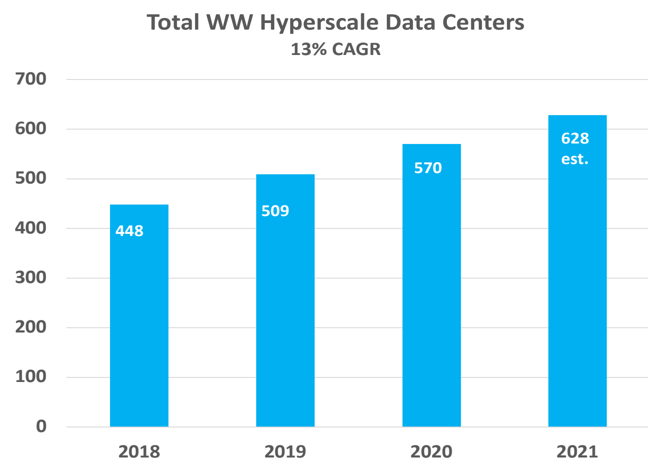

YE 2017	Hyperscale Profile	YE 2020 (est.)
386	Total WW HSDCs	~570
21%	Of all data center servers	47%
39%	Of all data center processing power	68%
34%	Of all data center traffic	53%
49%	Of all data stored in all data centers	57%

Source: Cisco Global Cloud Index, 2016-2021
https://www.te.com/usa-en/industries/data-center/insights/impact-of-evolving-data-centers.html

Hyperscale Data Centers Change the Rules

For companies who operate HSDCs, the cost may be the major barrier to entry, but ultimately that isn't the biggest issue – operations automation is. The rules of data center operations are different for HSDCs. They employ fewer tech experts than traditional data centers because they have used technology to automate so much of the overall management process. HSDCs have dedicated teams to initialize, configure, deploy software, and perform non-disruptive hardware maintenance. HSDCs abandon the operations constructs typically seen in conventional computing systems and favor stripped-down hardware designs that look to maximize the effectiveness of all available rack space. The sheer physical size of hyperscale data centers continually pushes existing architectural limits to new levels which presents constant challenges to their entire operation. The best example is the Citadel, a HSDC owned by Switch and is currently the largest data center campus in the world. The Citadel campus is located in the Tahoe Reno, Nevada business park and covers an area of 7.2 million square feet.

Pivotal Challenges Facing Hyperscale Data Center Operations

The efficient use of power, space, cooling and orchestrating dynamic, automated software management to minimize or eliminate downtime are crucial tenants for achieving ultra-efficient HSDC operations. Hyperscale carbon footprint, energy and sustainability strategies are providing valuable insights into what the future of computing might look like. Could it be possible for data center energy to function like an electric utility provider where you just plug in and pay for what you use someday?

Hyperscale Data Centers Reshape IT Landscape

- **A Hyperscale Data Center (HSDC) is an enormous distributed computing environment.**
- **Massive infrastructure - over 100,000 ft^2, largest campus is > 7.2 million ft^2 (= 88.9 soccer fields).**
- **Designed with self healing redundant components – if a failure, workload moves to another server.**
- **Automation and scalability are more critical than sheer size.**
- **Extreme energy consumption and carbon footprint challenges.**
- **Tape usage increasing and *will be critical* to enable HSDC growth and control infrastructure costs.**

Source: Horison, Inc.

Automation – HSDCs are using AI to automate troubleshooting instead of having simply alerting the IT department of the issue. This boosts the IT department's productivity as they will have more time to handle other IT issues. AI will also automate the ability to analyze nearly infinite permutations of co-dependencies that exist between an organization's physical and virtual infrastructure layers giving IT personnel a quick solution for how to resolve the issue before it impacts the end user.

DIY at Hyperscale Levels - HSDCs build their own custom systems from commodity components integrated into racks and supplied by Original Design Manufacturer (ODM) vendors. ODM products are restricted to a predetermined design compared to OEM products which can be made according to any specifications. Custom built, stripped down servers can be preloaded and shipped directly to the data center location decreasing installation time. HSDCs network servers horizontally by adding identical nodes to the existing cluster, enabling them to quickly be added or removed as capacity demands increase and decrease. It's imperative that *HSDCs try to optimize everything.*

Minimal Footprint and Dense Racks are Keys – Hyperscale computing abandons the high-grade constructs typically seen in conventional computing systems and favors stripped-down designs to maximize hardware efficiency. Rack management is a critical discipline for HSDCs and many CSPs. HSDC racks are typically wider than the standard 19" wide x 36" deep rack to allow for customization and usually range from 42U (73.5") to 48U (84"0) in height with 30+ units in a rack. Racks are getting taller and a 70U (122.5") open frame rack that pushes the boundary of rack mounting height is available. Note: one rack unit (U) is 1.75". Rack density is increasing as HSDC'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 is at a premium for HSDCs and *every cubic inch counts*.

Accelerated Tape Storage Adoption - Typically, 60-80% or more of all data is either infrequently or never accessed after 90 days. Traditional storage management techniques have often left large-scale data centers struggling with as much as 60% or more of data their data stored on the *wrong tier* of storage wasting millions of dollars a year. As a result, automated tape library usage in HSDCs is on the rise to easily scale capacity and to manage and contain out of control HDD growth by storing archival and unstructured data. Tape provides the lowest TCO, carbon footprint and energy costs and provides an additional level of cybercrime security via the tape air gap. As a result, several CSPs now offer specific archival and cold storage services based exclusively on tape.

Fast Scaling Bulk Storage Capacity – HSDCs 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. Tape can scale capacity much easier by simply adding media without the requirement for adding more energy demand. For HSDCs *computing and storage must scale non-disruptively and independently of one another.*

Intelligent Software Everywhere - Hyperscale storage management is heavily software-defined and is benefitting from AI and ML delivering a higher degree of automation and self-healing while minimizing direct human involvement. Load balancing software monitors the amount of data that needs to be processed, handles requests, and distributes resources to the available resources. When a server fails, workloads can be automatically moved from server to server or even data center to data center in some cases, avoiding repair or recovery actions and downtime as in a traditional data center. HSDCs implement high degrees of server virtualization to control server sprawl, with as many operating system images running on each physical server as possible. AI will also improve automated data migration between storage tiers to move low activity or inactive data from costly HDDs to more cost-effective tape. Advanced intelligent software is being used to manage the required high availability storage constructs including RAID, replication and erasure coding.

Security Challenges - When a HSDC or CSP owns the infrastructure, it's impossible to know who can access the network, who is permitted in the data center and whether security policies are actually enforced. Therefore, smart organizations encrypt their data before sending it, and they use their own encryption keys instead of the service provider's keys. For HSDC data stored on tape, customers have the added security advantages of encryption, WORM and the tape air gap to defend against cyberattacks.

Flash SSD challenges - Flash and solid-state memories play a major role in accelerating HSDC performance use less power and are easier to manage than HDDs. The Flash Translation Layer (FTL) intelligently manages everything from caching to performance, wear leveling, and to garbage collection, etc. The FTL overhead will be tolerated as the overall performance gains compared to HDDs are compelling.

Disk Drive Challenges - Once the storage infrastructure reaches hyperscale levels, traditional RAID and replication recovery architectures can become too expensive, too slow and unmanageable. The higher the disk capacity, the longer the RAID rebuild time required to restore or recover a failed drive to regain redundancy. A 4 TB disk will take at least 10 hours to rebuild – a 20 TB disk can take several days or even weeks which are unacceptable for HSDCs. With RAID reaching its practical limits in HSDCs, except for high 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 using a single namespace.

Hardware Upgrade and Refresh Cycle Frequency – The frequency of hardware upgrades can become labor-intensive and presents a key challenge for HSDC operations. The HSDC 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 can 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.

Cooling Technology is Critical for Hyperscale and CSP Energy Containment

Most data centers' largest operational expense is cooling their systems. Higher power consumption and increasing carbon emissions have forced HSDCs to develop new energy sources to reduce and more effectively manage energy. Managing cooling efficiency for racks is a critical discipline for HSDCs and CSPs. During the past twenty years, HSDCs have seen rack power density go up in parallel with compute and storage densities as more servers and more hard drives are placed into a single rack today than ever before. A typical IT rack used to consume one to three kilowatts, today it is possible to find loads of 20 to 40 kilowatts in a rack.

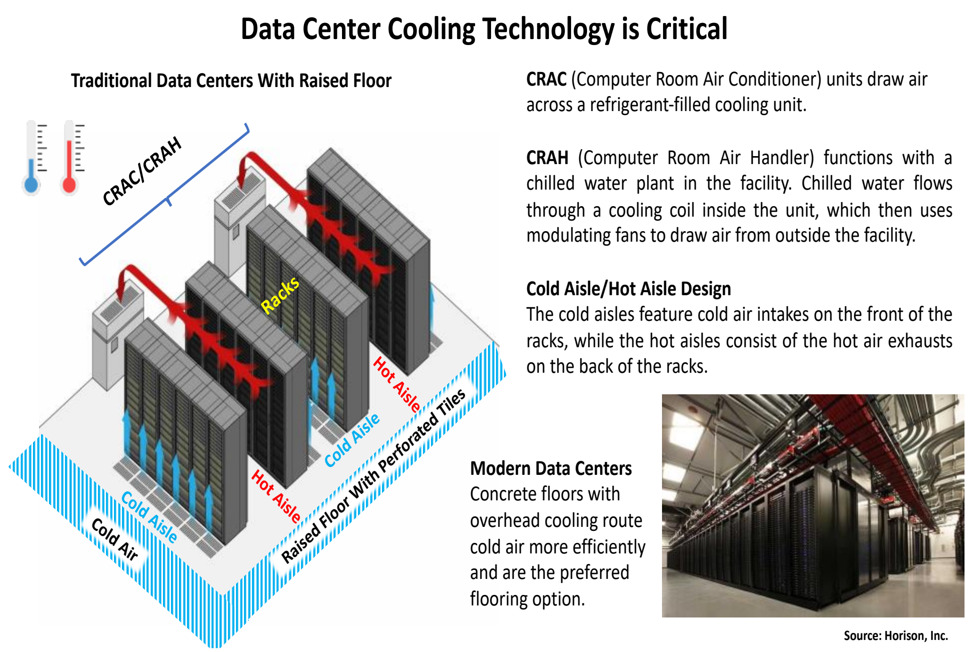

It used to be that data centers were simply built with raised flooring. This was done to allow for airflow, power and other cable routing, and flexibility during moves and changes. However, for most modern and larger scale data centers, concrete floors, or slabs, have become the more common choice, with cables and cooling systems running above the computer equipment rather than below offering a better job of routing cold air. Tape libraries use much less power and maximizes air handling and heat dissipation. Large scale data centers often allocate electrical power using a set number of kilowatts and fractions of megawatts. This enables data-center tenants to lease space in kilowatts, rather than the number of racks or square footage they might otherwise require.

Hyperscale PUE Pushing Maximum Efficiency

Electricity is the lifeblood of the IT industry. Reducing the number of servers and moving low-activity data from disk to tape present the greatest HSDC energy savings opportunities. High-density, multi-core data center servers typically use between 500 and 1,200 watts while HDDs use about 6-15 watts per hour, approximately three times more than flash SSDs. A typical desktop computer uses between 65 and 250 watts per hour. The goal for an efficient datacenter is to spend more energy on running the IT equipment than cooling the environment where the equipment resides. Over the past 10 years, overall data center power usage effectiveness (PUE), the ratio of total power required to run an entire facility versus the direct power involved in compute and storage has actually decreased. HSDCs are taking energy management to the next level by implementing Artificial Intelligence (AI), and Machine Learning (ML) to create a more precise PUE prediction model to anticipate and dynamically manage their energy consumption. Hyperscale operator Google now maintains an impressive PUE average of 1.11 across all its data centers which is very close to the *theoretically perfect* PUE of 1.0. Now that HSDC PUEs have drastically improved, the low-hanging fruit is gone. Even with HSDC efficiency gains, electricity demand is relentless presenting a number of challenges for the existing power grid and for new power sources.

Maximizing Hyperscale Energy Efficiency

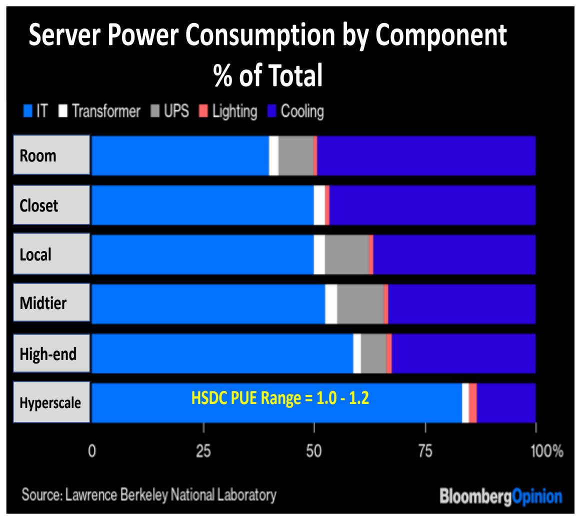

PUE	Data Center Power Efficiency Level
3.0	Very Inefficient
2.5	Inefficient
2.0	Average
1.5	Efficient
1.2	Very Efficient
1.0	Max Efficiency

Power Usage Effectiveness

PUE = Total Facility Power / IT Equipment Power

Avg. Data Center PUE = 1.67

AI and ML play key energy management roles.

The Value of Tape Rises Rapidly as Hyperscale Data Centers Grow

Today HSDCs are leveraging the many advantages of tape technology solutions to manage extreme data growth and long-term retention challenges. Some data requires secure, long-term storage solutions for regulatory reasons or due to the potential value that the data can provide through interpretation or analysis at a later date. Modern tape architectures allow HSDCs to achieve their data preservation and sustainability 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.

Carbon Footprint, Energy and TCO are Much 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. The substantial electricity use of data centers concern increases carbon dioxide (CO_2) emissions. HSDC and CSP migrations can [reduce CO_2 emissions](#) by as much as 60 million tons a year, according to new research from Accenture. 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, especially hyperscale, force them to continuously explore new cooling techniques. Studies indicate tape can provide a huge advantage for hyperscale TCO and carbon footprint reduction (by 86% and 87% respectively - see charts below). 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.*

Carbon Footprint and TCO Reduction Using Tape

Key Tape Advantages

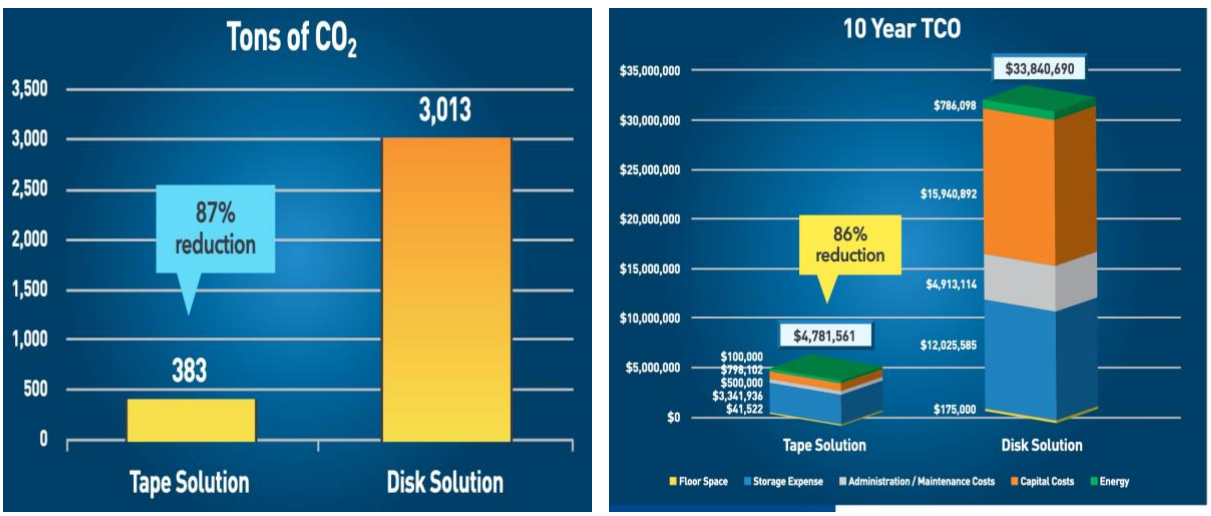

Ex: Assume 10 PB of cold data on disk growing at 35% per year for 10 years.
Use the publicly available Total Cost of Ownership (TCO) tool from the LTO Consortium.
Storing cold data on tape storage produces 87% less carbon footprint than disk.
Tape has a TCO that is 86% lower than disk for storing 10 PB of cold data.

Source: Brad Johns Consulting Reducing Data Center Energy Consumption and Carbon Emissions with Modern Tape Storage.

Tape Laboratory Demonstrations Project Few Capacity Limits

In December 2017, IBM and Sony demonstrated 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 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 storage media. Tape has a steeper areal density growth rate, currently at 34 % a year, than a disk drive which has a forecast growth of 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.

The Tape Air Gap Provides Security and Cybercrime Prevention Hyperscale Archives

Fighting the cybercrime epidemic has become a major focus for most data centers and HSDCs are no exception. Air gapped 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 removeable tape cartridge therefore preventing a malware attack on data stored on tape. HDD and SDD 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. Ransomware is projected to cost organizations globally as much as $20 billion by 2021. Air gapping should be an integral part of any archive, backup and recovery strategy whether on-premises or in the cloud.

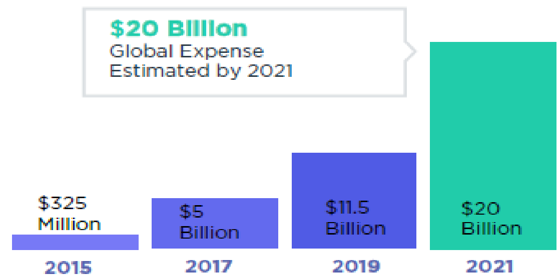

Global Ransomware Expense

Tiered Storage Becomes Strategic for Hyperscale Efficiency

HSDCs are taking advantage of tiered storage by more closely integrating high-performance SSDs, HDD arrays with automated tape libraries. Even though HSDCs, CSPs and enterprises are struggling with the exploding growth of disk farms which are devouring IT budgets and overcrowding data centers, many continue to maintain expensive disks often half full of data which 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 as its scalability, lower price, lower TCO and reduced carbon footprint plays a steadily increasing role as the size of the storage environment increases. The zettabyte era will fuel growth of enterprise, CSPs and HSDCs and will sooner or later make tape mandatory for sheer economic survival. For the hyperscale world *adding disk is tactical – adding tape is now strategic.*

Removable Media Plays a New Role for Moving Large Amounts of Data Quickly

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 moving tape media via a truck or airplane. Since tape media is removable (portable), tape becomes advantageous providing leverage if the cloud service provider shuts down or should you want or need to "quickly" move your entire petabyte media archive to another provider or 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 56 LTO-9, 18 TB cartridges which can 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.

Tape Enables the Most Cost-effective Object Storage Solution

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 and that trend looks to continue. 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 file or object format.

Growth of semi-structured and unstructured data has become one of the biggest infrastructure challenges for enterprise and hyperscale data centers as object storage requirements can quickly scale to hundreds of petabytes in a single namespace. Object storage software solutions using LTO and 3592 enterprise tape has emerged to provide a more economical solution to this challenge as an option for traditional file and block storage systems based on costly HDD and flash technology. Most unstructured data is infrequently used or inactive (cold), but still needs to be preserved indefinitely for future business analysis (big data) or compliance reasons making object data well suited for tape as a relief valve for large amounts of low-activity object data sitting on HDDs. By 2024 large enterprises are projected to triple their amount of unstructured data stored on-premises, at the edge and in the public cloud.

Modern Tape Value Proposition – 2021

Function	Benefits Summary
Price/TCO/CO_2	**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 Libraries are Available. Lab Demos Demonstrate Tape Capacities Can Reach 580 TBs.**
Open Standards	**LTO Tape Drives and LTFS Provide Open Standard Interface and Access.**
Scalability	**Tape Scales Capacity by Adding Media *Without Adding* Energy Consumption, HDDs Add Capacity by Adding Drives *With Adding* Energy Consumption.**
Energy /CO_2	**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	**No Foreseen Limits for Tape, HDDs Facing Areal Density and Performance Limits. Well Defined Roadmap.** Source: Horison, Information Strategies

Hyperscale Highlights

- The shift to Hyperscale is fueled by the migration of many smaller data centers to fewer, but much larger and more efficient CSPs and HSDCs.
- Large-scale data centers and many CSPs are becoming the next wave of hyperscalers.
- Energy, carbon footprint and sustainability are some of the most critical HSDC issues.
- Containing petascale and exascale storage environments on HDDs is becoming prohibitive.
- As the zettabyte era accelerates, tape will become mandatory for sheer HSDC economic survival.
- Remember that tape scales by adding more media and HDDs scale by adding more drives – and more energy.
- Given TCO and sustainability concerns, adding disk is now tactical – adding tape is now strategic.

Conclusion

The Covid-19 pandemic has forced our hand making digital transformation the name of the game in the IT industry. Eras are escorted out as others are ushered in, and the zettabyte era has fueled worldwide growth of HSDCs, CSPs and large data centers while further disrupting the traditional data center model. The HSDC represents the fastest growing data center segment today and has pushed carbon footprint, energy demand, and sustainability concerns to center stage. Today's HSDCs are re-engineering their storage strategies to cost-effectively manage extreme data growth and are now ready to take advantage the compelling economics and advantages of modern tape at scale. Without a new disruptive technology on the horizon to contain 60-80% of the world's archival and cold data, the many rich tape technology improvements of the past 10 years and the promising roadmap should feature tape as the most cost-effective storage solution for the unprecedented HSDC storage growth ahead. With over 50% of all data center data stored in HSDCs, tape will become instrumental in containing relentless capacity growth going forward. Just as it took vaccines, masks, hand washing and social distancing to slow down the growing wave of Covid-19 cases, it will take modern tape technology to slow down the growing wave of data center carbon footprint and energy consumption.