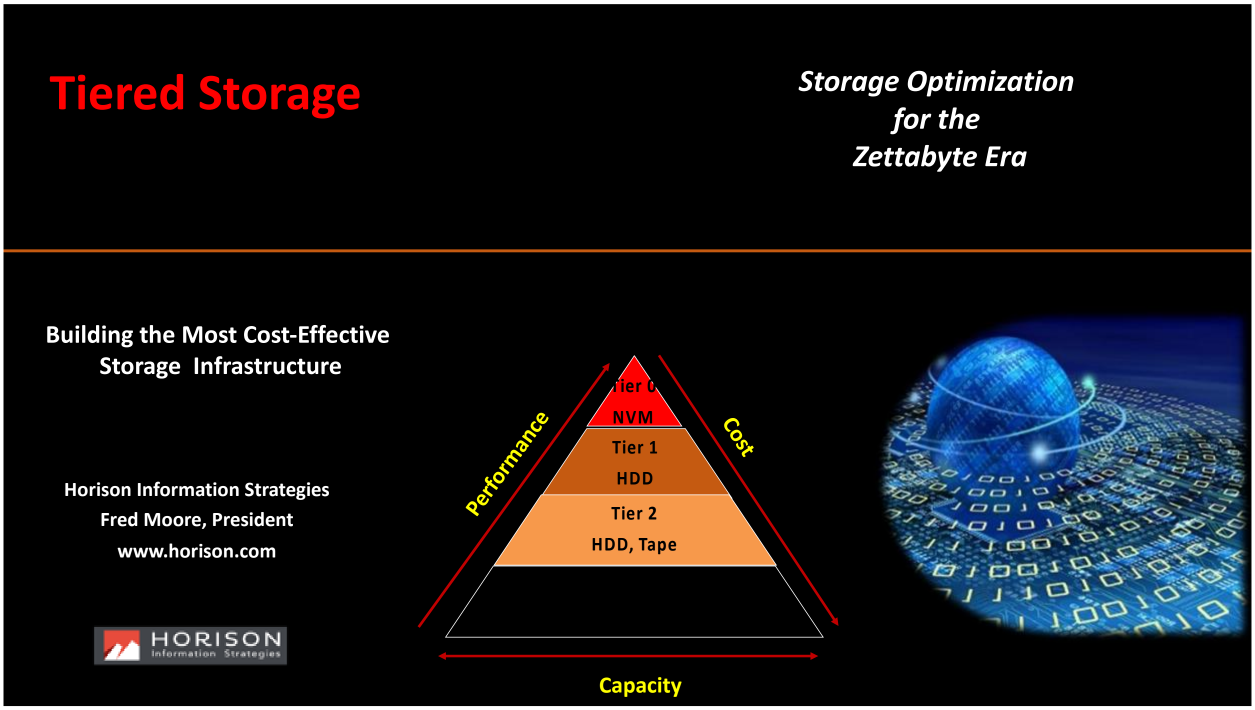

Introduction

The Covid-19 pandemic continues to affect everything from individuals to the global economy. Some businesses that primarily derived their value through in-person experiences weren't able to transform all of their offerings into digital content and their overall IT demand was reduced. Other businesses became dependent on at home workforce suddenly creating additional demand for online video/audio streaming, downloading, and video conferencing. The enormous dependency on the Internet throughout the lock down has led to the generation of vast amounts of data to be managed by data center facilities around the world. In parallel, the unknown magnitude of the Internet of Things (IoT), the big data archive pile up, mobile computing, surveillance, social media, and the huge amount of Covid-19 related data have lead the charge in data generation which is expected to grow at over 25% annually for the foreseeable future.

This data has to be stored somewhere. Fortunately, nearly all of this data resides inn data centers on SSD, HDD and tape storage and these technologies have never been more robust. Tiered storage is the process of determining which of these storage devices to use for which application at a given point in time and is a balancing act making trade-offs between frequency of access (performance), cost, availability and capacity. Tiered storage integrates hardware and storage management software to provide a seamless automated operation for businesses to realize the huge economic and carbon footprint benefits available from optimized storage implementations. The foundations of tiered storage had their beginnings over 30 years ago when disk, the arrival of automated tape libraries and advanced policy-based data management software such as (HSM) combined to effectively migrate less-active data to less expensive storage devices. The base business case for implementing tiered storage is compelling and becomes increasingly so as the storage environment gets larger.

How Big is the Digital Universe?

Industry estimates vary but the amount of data projected to be stored in 2025 is projected to be ~7.5 ZB according to IDC's Data Age report where 2,400 enterprise decision makers were surveyed from a variety of industries. The report also projects that 46% of the world's stored data will reside in public cloud environments by 2025. What we do know is that approximately 1.1 ZB of total storage capacity was shipped in 2020 across Non-Volatile Memory devices (SSDs), HDDs, and magnetic tape media, with HDDs capturing the majority (over 80%) of the total shipped capacity.

The Digital Universe
Archival Pileup by 2025 Could Reach 4.5 - 6.0 ZB

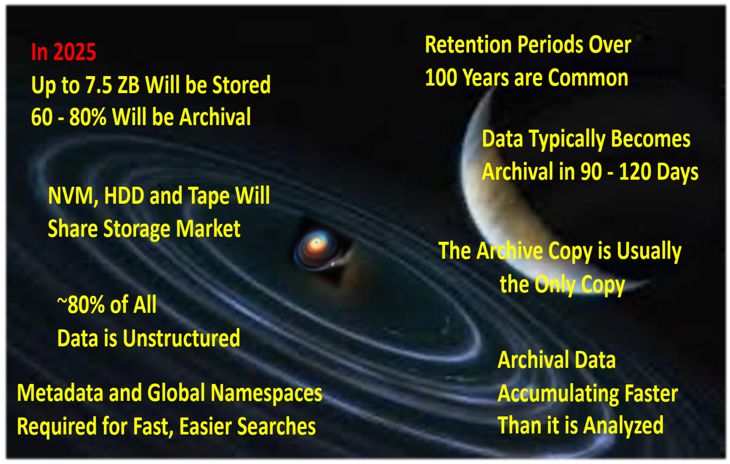

Today 60% or more of all digital data can be classified as archival and it could reach 80% or ~6 ZB by 2025, making it by far the largest and fastest growing storage class while presenting *the next great storage challenge*. Most archival data have never been monetized as its value remains unknown, but companies are quickly realizing that digital archives have great potential value. Companies looking to be relevant between now and 2025 will need to understand the role archive data can play in their organization's long-term success and how data archiving strategies and secondary storage platforms will evolve during that period. Given this trajectory, the archival storage paradigm will need to reinvent itself. Even as CIOs struggle with the exploding growth of disk farms, which are devouring IT budgets and overcrowding data centers, many continue to maintain expensive HDDs often half full of archival data which has had little or no activity for years. This is unfortunate and extremely costly. Archival data is much more cost effective when stored on tape than energy hungry HDDs. Obviously, as data center storage requirements soar, few can afford to sustain this degree of inefficiency and the value of tiered storage becomes obvious.

Tiered Storage Provides the Optimal Storage Infrastructure

Tiered storage is the process of assigning data to various storage technologies with the objective of reducing the total cost of storage. A tiered storage architecture places data in the storage hierarchy according to its business value. If expense were no object, businesses would use SSDs for all their storage requirements because they offer very high performance and reliability with no moving parts. However, for most businesses, storage costs are very important, as IT departments seek to live within their budgets and organizations strive to simultaneously minimize costs and maximize efficiency. Today's tiered storage hierarchy offers a range of robust technologies ranging from ultra-high capacity, low-cost storage at one end of the hierarchy to very high levels of performance and functionality and at the other making a highly effective storage implementation a reality.

Inside the Storage Tiers – Examining the Technologies

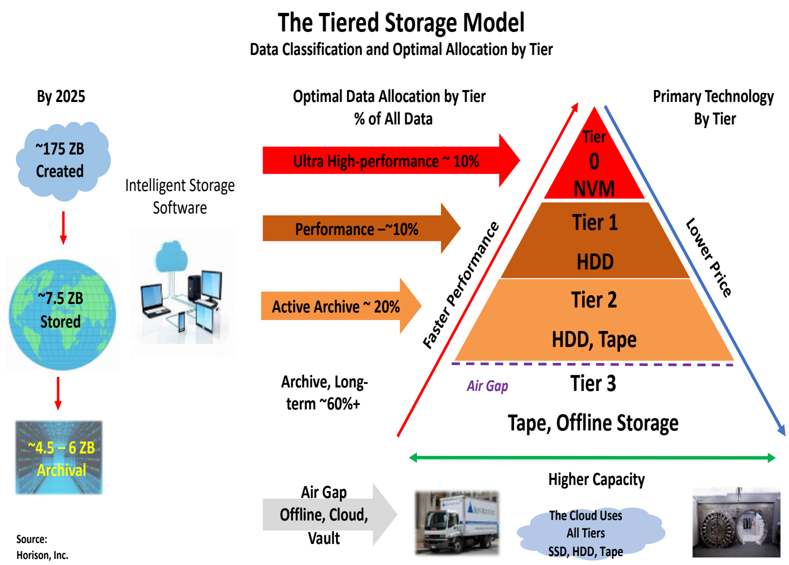

Today's de-facto standard storage hierarchy is defined by four different technology tiers. The minimal tiered system has two tiers, one for frequently accessed data and one for less frequently used data. Deploying just two storage tiers offer a limited choice over where a given file should be stored. The more tiers that are available, the more choices administrators have regarding the placement of data, and the greater the storage cost savings. Tiers are delineated by five primary attributes: price, performance, capacity, availability and throughput. As storage pools grow, an automated tiered storage environment becomes the optimal, most cost-effective storage architecture since 1) manual data movement is complex and time consuming 2) the amount of digital data is continually increasing 3) most data is stored in the wrong storage tier, and 4) limited staff resources often leave storage administrators stretched too thin. The CSP (Cloud Service Provider) and offline data vaults represent storage services which can use any or all the four technology tiers. Leading CSPs provide special archival services with retrieval times ranging from a few minutes to several hours, offering encrypted storage in geographically distributed locations for higher availability. Many of these services are based on tape.

Leading storage suppliers offer a complete tiered storage portfolio including high-performance NVM (Non-volatile Memory) SSDs (Solid State Disk), HDD (hard disk drive) arrays and automated tape libraries to address data requirements throughout the data lifecycle. SSDs remain expensive compared to HDDs though the price gap is closing. That means that SSDs need to be used judiciously for storing high performance data. Data that is less critical to day-to-day operations can be stored on lower cost HDDs and data that is less frequently accessed, retained for legal requirements and big data archives waiting to be analyzed is best stored on air-gapped tape systems. The greatest benefits of tiered storage occur when archival data is migrated to tier 3 (the tape tier) since this tier is the most cost-effective home for 60% or more of all archival and cold data and it offers the greatest differential in acquisition price, TCO and carbon footprint compared to tiers 0, 1 and 2.

The Tiered Storage chart below serves as a reference model to describe the key technology attributes of the four de-facto standard storage tiers.

Inside the Storage Tiers - The Physical View

Primary Storage → Secondary Storage

Storage Tier	Tier 0 Ultra-high Performance	Tier 1 Performance	Tier 2 Active Archive	Tier 3 Archive, Long-term
Amount of Digital Data in Each Tier	~10%	~10%	~20%	60% (or more archival data)
Primary Technology	NVM (DRAM, NAND Flash, PCM, 3D-Xpoint)	Enterprise disk arrays FC, SAS	Midrange disk arrays SATA	Tape libraries, offsite data vaults, cloud services
Nominal Access Time	1-10 µ	5-10 ms	5-20 ms	25-120 sec
Data Transfer Rates	550/520 MB/s R/W	160-260 MB/s	80-260 MB/s	400 MB/s LTO and Enterprise
Typical File Access	Random/Seq.	Random/Seq.	Random/Seq.	Sequential only
Data Classification Category	I/O intensive, response-time critical, OLTP, ultra high-performance	Mission-critical, OLTP, revenue generating applications	Vital, sensitive, business important applications	Archives, fixed content, big data, reference data, govt. regs, high data rates
2021 est. price *	~$120/TB	~$35/TB	~$20/TB	~$4-5/TB
Reliability (BER)	$1x10^{17}$	$1x10^{16}$	$1x10^{15}$	$1x10^{20}$
Media Life	3-5 years	4-5 years	4-5 years	> 30 years
Power Rating	3-5 W	6-15 W	6-15 W	Lowest

Source: Horison Information Strategies

***Prices are relative and vary based on a variety of factors. See specific vendor for current quotes.**

Tier 0: Ultra High-performance Storage - Tier 0 is uses solid state memory technologies to store high performance, high IOPs data that needs to be captured, analyzed, and retrieved at the highest possible speed. The emergence of Storage Class Memories including the popular 3D NAND Flash, and newer technologies such as Phase Change Memory (PCM) and 3D XPoint, are expected to expand tier 0 solutions as the future roadmap for NVM solutions is robust. Tier 0 solutions are the most expensive tier on a $/TB basis though cost is typically not the major factor for selection. HDDs have access times in milliseconds, while SSDs operate in microseconds and generally require half to a third of the power of HDDs. Flash SSDs have surpassed HDD technology in areal density with announcements of 2,770 Gb/in^2. With the transition from single layer to 3D-NAND flash offering 64-96 layers and 160-layer Flash under development, the SSD industry continues to take market share from the HDD industry as the price differential between SSD and HDD shrinks. The NVMe and new NVMeOF (NVMe Over Fabric) interface was designed from the ground up for SSDs and is now the primary SSD interface specification for new storage systems, replacing the disk-based SATA, SCSI, FC and SAS interfaces.

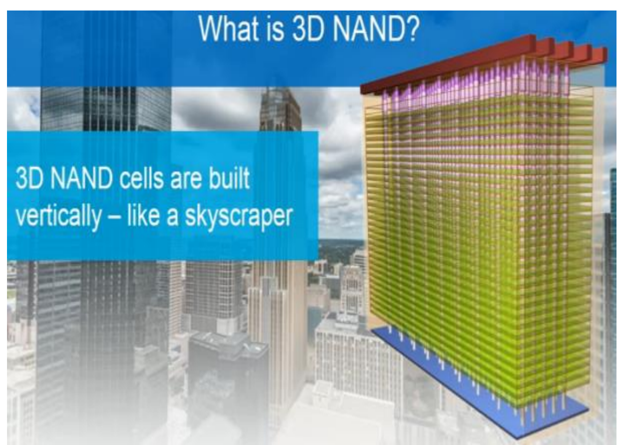

Tier 1: Performance Storage - Tier 1 is used for mission critical data and uses enterprise-class HDD systems requiring good performance, high availability with near-zero downtime and fast RTOs (Recovery Time Objectives) to support customer-facing and revenue-generating applications. HDDs often use mirrored or double parity RAID for fast recovery of mission critical data. HDDs are facing scalability (access density) challenges as capacity continues to increase with few if any corresponding performance (IOPs) improvements. Looking ahead, HDDs face some challenges for increasing capacity, though HAMR and MAMR developments should help provide sustainable capacity growth for standard 3.5-inch nearline drives. To increase capacity, some HDD suppliers have increased the number of platters from three to seven while using helium filled disk enclosures to reduce friction. To further improve access time, dual actuator drive technology arrived in 2020. By using two read/write heads, a disk drive is divided into two logical halves that perform read/write operations concurrently to improve IOPs functioning as two HDDs in one enclosure. Four actuators have been considered. As manufacturers push the evolution of HDD technology by increasing platters and adding new features such as more platters and multiple actuators, the total parts count, cost and complexity continue to increase, and the drives can become less stable. The typical HDD will stay in service 4 – 5 years before replacement.

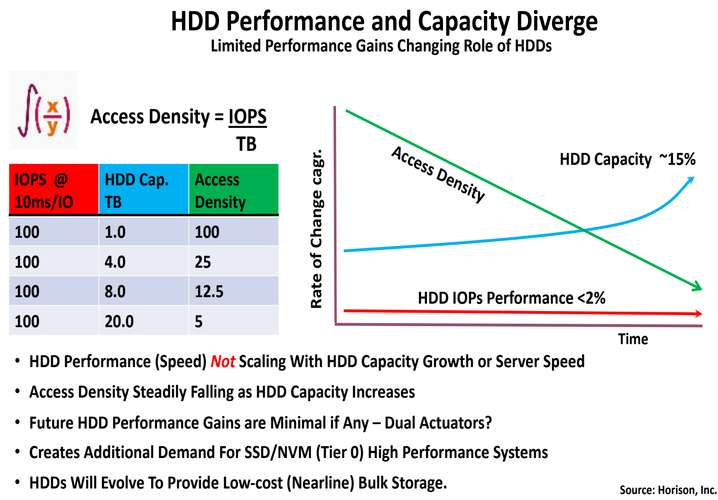

Tier 2: Active Archive Storage - This tier includes high capacity nearline HDD storage addressing less time-critical data at lower price points than tier 1 drives. Tier 2 supports a broad range of major business applications including low-activity databases, batch workloads, active archives (nearline disk and tape combined), backup, email, file systems and ERP (Enterprise Resource Planning). Tier 2 solutions must safely store active business data where sub-second response and the highest availability are not a requirement, but reasonable response time still is needed. Choosing the optimal tier 2 solutions is normally a balance between cost and performance.

Tier 3: Archive, Long-term Storage - Tier 3 represents the largest potential storage tier as it is the optimal storage system for at least 60% or more of all stored digital data. As most data ages, access activity drops off rapidly and data typically reaches archival status in 90 - 120 days becoming cold data. Low cost is the overriding decision factor for tier 3 storage. Fortunately, modern tape has a media life of 30 years or more making it best suited to address long-term data retention requirements. New tape solutions are arriving allowing for objects and metadata to be efficiently written and read to and from tape in native form. Tape is presently the most cost-effective choice for tier 3 data given it has the lowest $/TB, TCO and carbon footprint of any storage option. Steady advances point to tape continuing to deliver the lowest cost, highest capacity, fastest data transfer rates, the most reliable digital storage medium available, with reliability levels three orders of magnitude better than the best HDDs.

Looking ahead, the tape industry has pushed capacity, reliability and media life to record levels surpassing disk drives. In April 2015, IBM and Fujifilm demonstrated a 220 TB (raw) tape with 123 gb/in^2 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/in^2. In Dec. 2020, IBM and Fujifilm demonstrated a record areal density of 317 gb/in^2 yielding a 580 TB cartridge using a new magnetic particle called Strontium Ferrite (SrFe). Laboratory demonstrations indicate steady advancements with few limitations in tape technology for the decade ahead and may signal the way for a new, very low-cost mass storage tier by 2025. See technology innovation chart above.

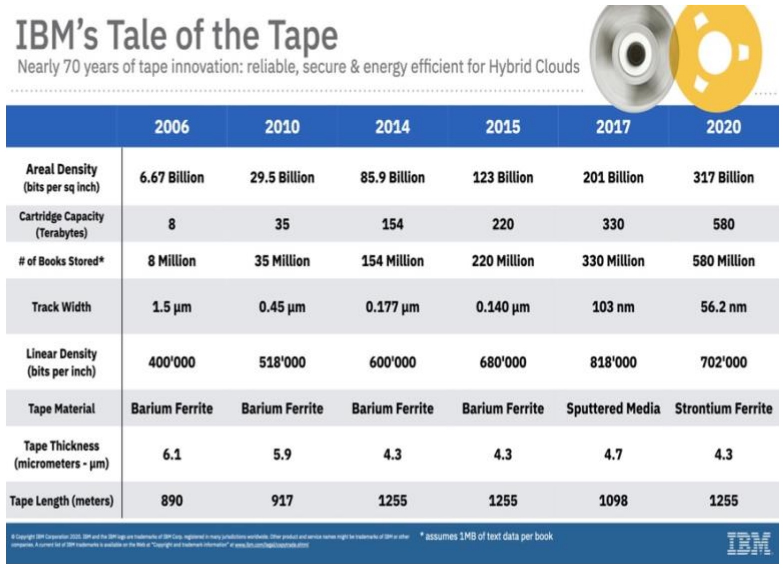

	2006	2010	2014	2015	2017	2020
Areal Density (bits per sq inch)	6.67 Billion	29.5 Billion	85.9 Billion	123 Billion	201 Billion	317 Billion
Cartridge Capacity (Terabytes)	8	35	154	220	330	580
# of Books Stored*	8 Million	35 Million	154 Million	220 Million	330 Million	580 Million
Track Width	1.5 µm	0.45 µm	0.177 µm	0.140 µm	103 nm	56.2 nm
Linear Density (bits per inch)	400'000	518'000	600'000	680'000	818'000	702'000
Tape Material	Barium Ferrite	Barium Ferrite	Barium Ferrite	Barium Ferrite	Sputtered Media	Strontium Ferrite
Tape Thickness (micrometers - µm)	6.1	5.9	4.3	4.3	4.7	4.3
Tape Length (meters)	890	917	1255	1255	1098	1255

* assumes 1MB of text data per book

A New Storage Tier by 2025?

At some point, today's storage hierarchy will need to evolve providing a highly scalable and ultra- capacity solution to address the enormous tier 3 archival avalanche. The potential emergence of a *Deep Archive* (cold) storage tier will need to provide highly secure, long-term storage and require minimal remastering cycles. Remastering is a labor-intensive process that typically migrates data to new media every ~5 - 10 years. Deep archives will offer "infinite" capacity scaling and media life of 100+ years. The optimal design will have a robust roadmap, minimal environmental/carbon footprint, a price below all other competitive offerings. Fast access times will not be the critical factor for a Deep Archive tier – lowest cost will be.

Data Lifecycle Behavior Patterns

Archival data will continue to be the largest and fastest growing data classification segment. For most data types, the probability P(a) (probability of access) begins to fall one month after creation and typically falls below 1%, between 90 - 120 days from creation. Some data becomes archival upon creation and can wait years for reference or further analysis adding to the archival pile-up. Long lifecycles are becoming common as retention periods can exceed 100 years while some data will never be deleted.

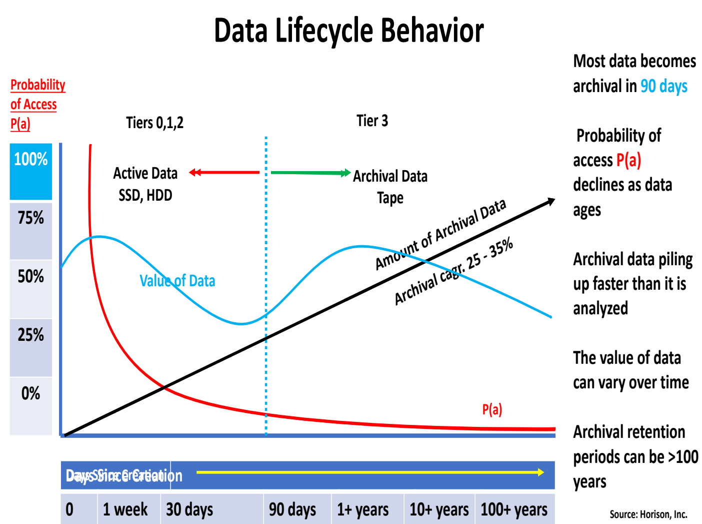

Once data reaches pre-defined thresholds, intelligent software moves the data to the optimal tier. If the access frequency has increased, it is moved up to a higher performing tier. If data has not being used over a period of time, it is moved down to a lower-cost, higher-latency tier. Three consistently observable profiles have evolved that improve understanding data behavior over its lifecycle. See chart above:

1) **the probability of access P(a)** for most data declines as the data ages
2) **the value of specific data** to a business can change over time based on a variety of circumstances
3) **the amount of data** steadily increases with age since more data is being kept for much longer periods of time than ever before pushing older and less active data into tier 3 status

Data Classification – Understanding Your Data Access and I/O Patterns

The process of organizing and storing data by relevant categories so that it may be used and protected more efficiently is called data classification. A key aspect of a storage tiering architecture is how to classify data into levels of importance and assign it to the appropriate storage tiers. Data classification is of particular importance when it comes to risk management, compliance, and access security. Security levels are described in the de-facto standard Information Security Scale.

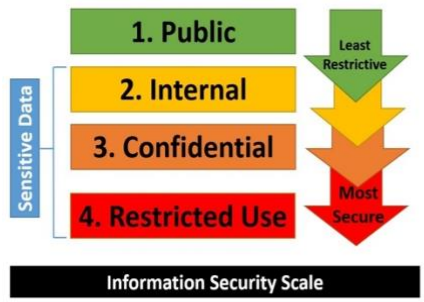

Data is typically classified into four general categories which are mapped to the storage tier best fitting those characteristics. Classifying data and understanding the access characteristics and reference patterns of applications has been a time-consuming and difficult task. The data classification process must be smart enough to enable rapid classification of large volumes of data. With an understanding of performance, access patterns and security, you can map the storage solution that best suits the application requirements. Learning about or characterizing data access patterns, while extremely important, is not enough. The characteristics of the storage infrastructure itself, both hardware and software, must be understood as well. Look for AI to significantly improve this process in the near future.

Data Classification by Tier – Getting Started

Data classification aligns data with the optimal storage tier. If effective data classification software tools are not readily available, the classification table below can serve as a simple starting point to begin the classification process by mapping the required data characteristics for a given application to the optimal tier. When classifying data, ask yourself:

1) What are the performance requirements (initial access time, IOPs) for this data?
2) When does this data reach archival status?
3) How long do I need to retain the data?
4) How will large data archives be managed?
5) How soon do I need the data back if lost, damaged or inaccessible (RTO)?
6) How secure does it need to be? (How critical is this data?)
7) What regulatory requirements need to be adhered to?

Tier and Storage Class	Key Applications by Tier
Tier 0 **Ultra-High Performance** **Mission Critical**	High performance databases, operating system files, OLTP, reservation systems, indices, logs, roll files, directories, system catalogues, HPC and scientific applications, real-time analytics/simulation, database acceleration, artificial intelligence and machine learning, augmented reality, any data that demands the highest levels of I/O performance (IOPs)
Tier 1 **Performance** **Mission & Business Critical**	Mission critical databases, tele-medicine, online financial systems, navigational systems for a spacecraft, reservation systems, ATM, Point of Sales, virtual machines, police, military and national security systems, railway/aircraft/transportation operating and control systems, electric grid and power systems, nuclear reactor controls
Tier 2 **Active Archive** **Warm Data**	Business-critical applications, surveillance, Internet applications, data protection – backup, recovery, security systems, image capture and retrieval, application development and test, data warehousing, ERP, big data, mobile devices, edge devices, BC/DR
Tier 3 **Archival** **Long-term Storage** **Cold Data** ***Deep Archive…***	All long-term data retention, archive and backup, big data – yet to be analyzed, compliance data, GDPR, medical records, photos and images, e-mail history, documents, objects, unstructured data, scientific, video, movies, audio, social media history, archive cloud applications, surveillance and security system history, off-site media storage, remote data vaults, BC/DR, CSP and HSDC archival services

Artificial Intelligence (AI) and Metadata to Improve Data Classification

If data is the new currency, then storage is the new bank. Today 50% of the top 10 companies are data-based platforms representing a fundamental shift in terms of the way the market views the value of data. IT staffs are under increasing strain as the volume and complexity of managing daily workloads defy the traditional approach of simply adding more HDDs when capacity is maxed out. Furthermore, data will have to be analyzed, tagged (using metadata), processed and subjected to other processes in order to effectively support analytics and the real-time applications that drive productivity. Organizations will need to dynamically migrate, recall, replicate, mirror or spread data across increasingly complex and geographically dispersed tiered infrastructures, no simple task.

Tiered storage and data management are poised to get a dose of intelligence – meaning AI (Artificial Intelligence). Artificial intelligence is the theory and development of computer systems to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision making, and data management in the near future. AI is poised to improve data management by gathering large amounts of information about how the data is used and building policies to manage the data. Who access it? How often is it accessed? What kind of file is it? What type of data (critical, non-critical) does it contain? By gathering this information, learning algorithms can start to create predictive models describing how the data will most likely be used in the future and to place data in the optimal tier.

Tape Fuels Tiered Storage Economics

The greatest benefits from tiered storage occur by moving as much data as possible to the tape tier (3). The more data that can move to tape (tier 3), the greater the carbon footprint and TCO savings over an all-disk approach. Most large-scale enterprises now have corporate sustainability goals. The study below assumes 10 PB of cold data on disk growing at 35% per year for 10 years. Using the publicly available [TCO Calculator]() from the LTO Consortium, storing cold data on tape storage produces 87% less carbon footprint than disk and has a TCO that is 86% lower than disk. Furthermore, the larger the storage pool becomes, the greater the economic tiered storage benefits become. This is particularly true for CSPs (Cloud Service Providers) and HSDCs (Hyperscale Data Centers) whose archive requirements can be from 10 to 100x greater. Surprisingly, most of the largest HSDCs are in hot or temperate climates, consuming vast amounts of energy to keep them from overheating. Almost as important as switching data centers to low-carbon energy sources is improving their energy efficiency. Much of this comes down to the energy needed to keep the processors cool. Of the world's 10 largest HSDCs, two are in the desert heat of Nevada, and others are in Georgia, Virginia, and Bangalore. The TCO chart below compares 10 PB of cold data on disk growing at 35% per year for 10 years to automated tape.

Carbon Footprint and TCO Reduction Using Tape

Key Tape Advantages

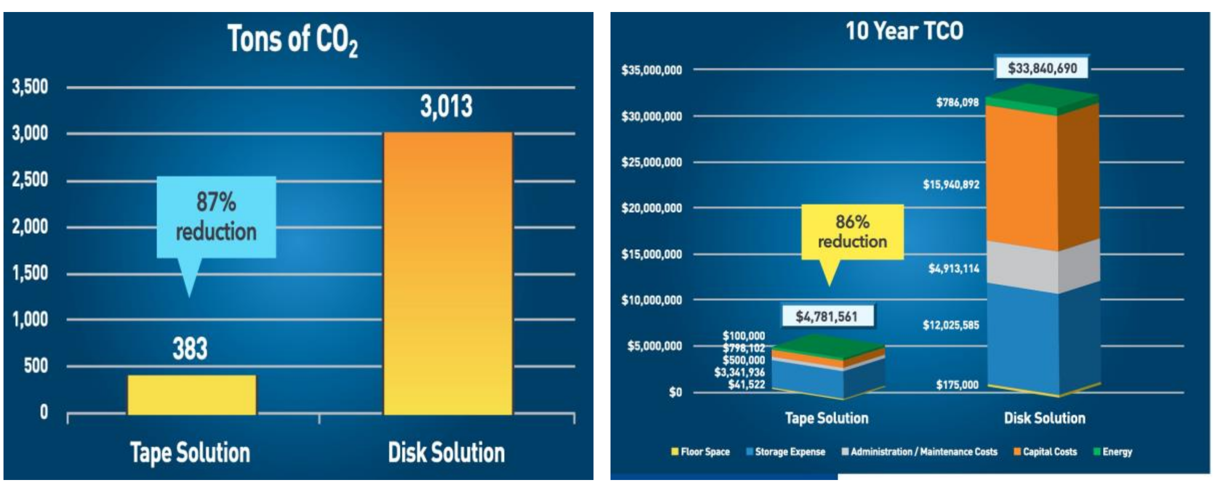

Ex: Assume 10 PB of cold data on disk growing at 35% per year for 10 years.
Use the publicly available Total Cost of Ownership (TCO) tool from the LTO Consortium.
Storing cold data on tape storage produces 87% less carbon footprint than disk.
Tape has a TCO that is 86% lower than disk for storing 10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

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

Key Benefits of Tiered Storage

Tiered storage provides several key benefits, the most important of which are:

Reduced storage costs: by storing each data class on the lowest cost storage that offers the required performance companies can avoid paying for higher-cost storage that isn't needed. This reduction in overall storage costs, TCO and carbon footprint is the primary driver for implementing a tiered storage system and the greatest savings occur when the tape tier (tier 3) is used.

Higher storage efficiency: High performance HDDs which rely on RAID or short stroking (formatting a disk drive such that data is written only to the outer sectors of the disk's platters to reduce seek time) are storage inefficient because the amount of data that can be stored on these systems is less than the total amount of storage capacity provided. Since tiered storage systems can relieve the demand on RAID and short-stroked storage by enabling data that does not require such high performance to be moved to lower storage tiers, storage efficiency will go up.

Ability to reuse older storage equipment: Storage tiering can provide a new lease of life for storage systems which would otherwise need to be replaced because they no longer offer sufficient performance for upper tier data uses. After tiering is implemented, they can often be used in a lower cost tier storage.

Conclusion

An effective tiered storage strategy has to balance storage costs, data lifecycle management practices, and storage technology capabilities. The best choice is to take advantage of a multi-tier storage system that automatically maps your data to the most cost-efficient (TCO) tiers of storage. The storage industry is anxious for AI management tools to simplify this process, but they won't suddenly appear - rather they will evolve and improve their capabilities over time. Maintaining large amounts of data on the wrong storage tier has costly consequences that will only mount as the storage pool grows. The long-standing storage management goal of delivering on the promise "to have the right data - in the right place - at the right time" can finally become a reality with tiered storage. For some enterprises, building a tiered storage strategy may still be optional but that day may abruptly end. With data storage volumes expected to continue to grow at over 25% annually and the unknown IoT storage demands from billions of data generating end points around the globe looming, the time to develop a tiered storage strategy may already be in your rear-view mirror.

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.