2020 Technology Update Series

Tiered Storage

Building the Optimal Storage Infrastructure

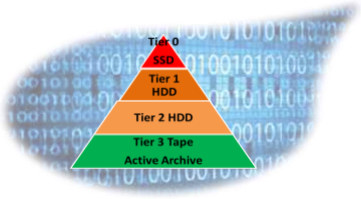

HORISON Information Strategies

Fred Moore
President
Horison Information Strategies
Horison.com

Introduction

Fortunately, as data continues to grow exponentially, the selection of data storage technologies has never been more robust. The choice of what storage device to use for which application at a given point in time is a balancing act making trade-offs between frequency of access (performance), cost, and capacity. Storage tiering has become a key strategy that lets you optimize the use of storage resources, save costs and make the best use of storage technology for each data classification. The foundations of tiered storage had their beginnings over 30 years ago when disk, 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. Tiered storage integrates hardware and storage management software to provide a seamless operation and for customers to realize the huge TCO and ROI economic benefits available from optimized storage implementations. The business case for implementing tiered storage is compelling and becomes increasingly so as storage pools get larger. Today's storage tiers offer several 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. The non-stop, increasing growth of data will require the continual evolution of new, more advanced approaches to tiered storage and management capabilities.

Tiered Storage – A Closer Look into the Storage Tiers

Today's de-facto standard storage hierarchy is defined by four tiers. The minimal storage tiering system has two tiers, one for frequently accessed data and one for archival data. The more tiers that are available, the more choices administrators have regarding the placement of data, and the more efficiently storage resources can be utilized. Tiers are delineated by differences in four primary attributes: price, performance, capacity and functionality.

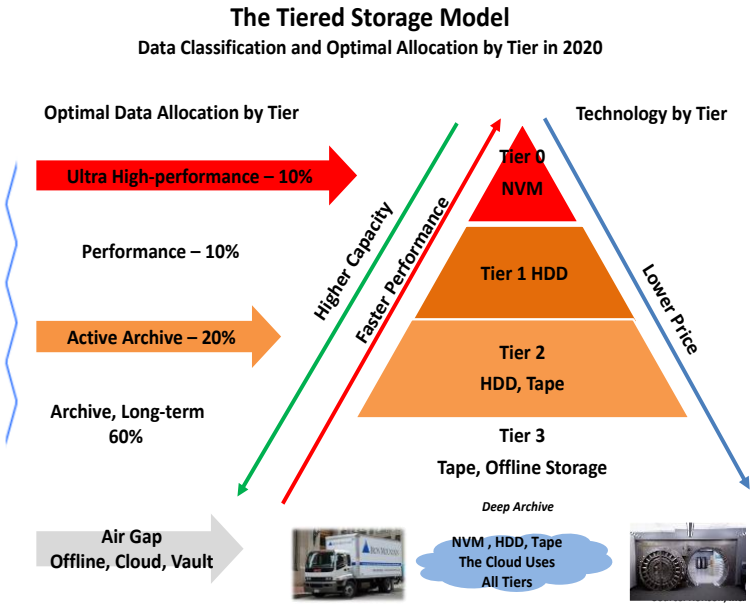

As storage pools grow, an automated tiered storage environment becomes the optimal, most cost-effective storage architecture since 1) manual data movement is complex and time consuming 2) the amount of digital data is continually increasing 3) most data is stored in the wrong storage tier, and 4) limited staff resources often leave storage administrators stretched too thin. The CSP (Cloud Service Provider) and offline data vault are storage options representing services which can use any or all of the four technology tiers.

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 its lifecycle. The benefits of tiered storage are greatest when the tape tier (tier 3) is used for archival data as it has a significantly lower acquisition price and TCO compared to other tiers.

The amount of data stored in 2025 is projected to be ~7.5 ZB according to IDC's 2019 "Tape and Cloud: Solving Storage Problems in the Zettabyte Era" white paper. At least 60% of all data can be classified as archival and it could reach 80% or more by 2025, making it by far the largest and fastest growing storage class while presenting *the* next great storage challenge. Future software and storage technology developments are on the horizon indicating substantial advancements in price, performance, reliability, capacity, security and throughput capabilities throughout the hierarchy. To further the tiered storage value proposition, the tape industry continues to innovate and deliver compelling new features with the lowest economics and the highest reliability levels. This has established tape as the most cost-effective choice for archiving and tier 3 as well as playing a larger role for backup, business resumption and disaster recovery.

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

Inside the Storage Tiers - The Physical View 2020

Storage Tier	Tier 0 Ultra-high Performance	Tier 1 Performance	Tier 2 Active Archive	Tier 3 Archive, Long-term
Amount of Data in Each Tier	10%	10%	20%	60% (or more archival data)
Primary Technology and Interface	NVM (DRAM, 3D-Flash SSD, PCM, 3D-Xpoint)	Enterprise disk arrays	Midrange disk arrays – scale out	Tape libraries, offsite data vaults, cloud services
Nominal Access Time	1-10 µ	5-10 ms	5-20 ms	25-121 sec
Data Transfer Rates	550/520 MB/s R/W speed 3,500/2,300 MB/s R/W speed	160-220 MB/s	80-220 MB/s	360 MB/s LTO 400 MB/s 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
2020 est. price *	~$140/TB	~$40/TB	~$30/TB	~$4/TB
Reliability (BER)	$1x10^{17}$	$1x10^{16}$	$1x10^{15}$	$1x10^{19}$
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

***Prices are relative and vary based on a variety of factors. Tape price includes 2.5x compression. See specific vendor for current quotes.**

Tier 0: Ultra High-performance Storage - Tier 0 is used to store extremely 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 (NVMe), (Non-Volatile Memory Express) and Phase Change Memory (PCM) and newer technologies such as 3D XPoint, are expected to expand tier 0 solutions as the future roadmap for NVM solutions is robust. The arrival of the NVMe specification that allows an SSD to make highly effective use of a high-speed PCIe bus in a server, has provided a huge boost for further increasing I/O performance. Tier 0 solutions are the most expensive tier on a $/GB basis though cost is typically not a major factor for selection. High performance flash SSDs generally require half to a third of the power of HDDs. Hard drives have access times in milliseconds, while SSDs operate in microseconds. 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-flash with 64-96 layers and potentially more layers, the SSD industry continues to take market share from the HDD industry.

Tier 1: Performance Storage - Tier 1 is used for mission critical data and uses enterprise-class disk systems requiring good performance, high availability with near-zero downtime and fast RTOs (Recovery Time Objectives) to support customer-facing and revenue-generating applications. When fully featured, these systems can carry lots of functionality at a premium $/GB price, but this is justified because lower performance or less reliable disk could directly impact customer satisfaction, business revenues and corporate viability. Mission critical data often uses mirrored or double-parity RAID for fast recovery purposes. Looking ahead, disk drives can expect to see minimal performance improvements and face substantial challenges for increasing capacity, though HAMR and MAMR may help provide sustainable capacity growth. Disk is facing scalability (access density) challenges as capacity continues to increase without any corresponding performance (IOPs) improvements. The typical disk will stay in service between 4 – 5 years before replacement.

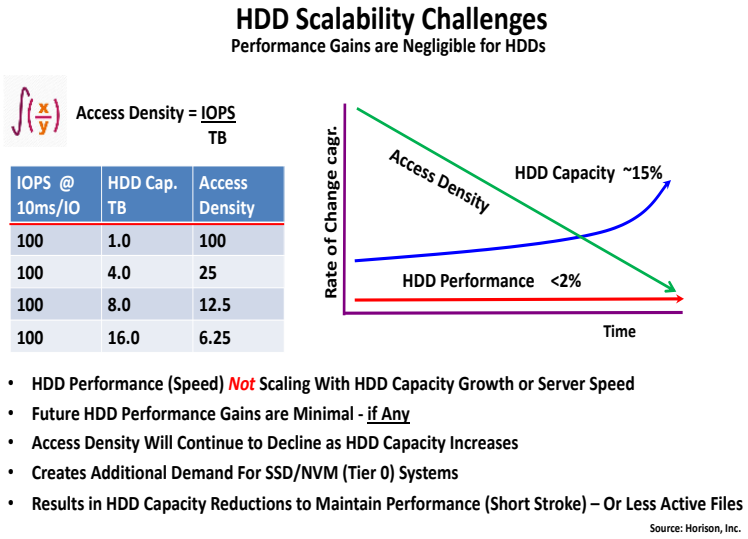

Tier 2: Active Archive Storage - This tier includes midrange disk storage addressing less time-critical data at lower price points than tier 1 disk. Tier 2 supports a broad range of major business applications including low-activity databases, active archives (nearline disk and tape combined), backup, email, file systems, batch workloads and ERP (Enterprise Resource Planning). Tier 2 solutions must securely store active business data where sub-second response is not necessarily 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 storage represents the archival segment which comprises at least 60% or more of all 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 $/GB and TCO of any storage option. The tape industry has pushed capacity, reliability and media life to record levels surpassing disk drives. Advanced laboratory demonstrations indicate steady advancements with few limitations in tape technology for the decade ahead while pushing cartridge capacities into the 300 – 400 TB range.

Data Classification – Understand and Learn Your Data Access and I/O Patterns

Data classification is key to tiered storage. In practice, data is typically classified into four general categories which are mapped to the storage tier best fitting those characteristics. Classifying data and understanding the I/O characteristics and reference patterns of applications has long been a tedious and often extremely difficult task. The data classification process must be smart enough to enable rapid classification of large volumes of data. With an understanding of I/O activity and access patterns, you can tailor the storage solution to the application. 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 improve this process in the near future.

Tiered storage can be effectively managed at a much lower cost by deploying advanced tools that move, migrate, protect, deploy global namespaces, add metadata and respond to changing lifecycle needs of the data. Three consistently observable profiles have evolved that improve understanding data behavior over its lifecycle.

The Data Lifecycle Model

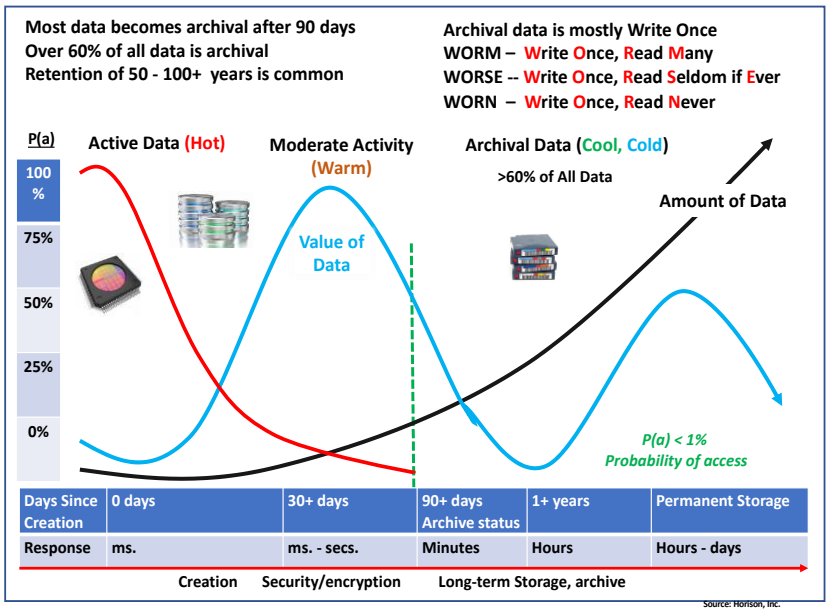

See chart above:

1) **The probability of reuse** of most data declines as the data ages
2) **The value of data** to a business can change over time based on a variety of circumstances
3) **The amount of data** increases as it ages 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

Since data has become the most valuable asset for many businesses, managing and protecting data throughout its lifetime has become their most critical storage management task. With or without sophisticated tools, most organizations should be able to identify their most critical applications and have a clear understanding of the value they provide to the success of the business. In the past, tiered storage implementations have often been limited due to:

1) Lack of automated data classification capability and limited deployment of data mover tools
2) Lack of support for heterogeneous server and/or storage environments (siloed systems)

A Closer Look at Data Classification by Tier

Data classification aligns data with the optimal storage tier. If data classification software and 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) How fast (the performance requirement) do I need to access the data?
2) How long do I need to retain data?
3) How will unstructured data be managed?
4) How soon do I need the data back if lost, damaged or inaccessible?
5) How secure does it need to be? (How critical is this data?)
6) What regulatory requirements need to be adhered to?

Classify Applications by Storage Class and Tier	
Ultra-High Performance	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, any data that demands the highest levels of I/O performance (IOPs), artificial intelligence and machine learning, augmented reality
Performance	Mission critical databases, tele-medicine, online financial systems, navigational system 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
Active Archive	Business-critical applications, Internet applications, data protection – backup, recovery, security systems, surveillance, image capture and retrieval, application development and test, data warehousing, ERP, big data, mobile devices, edge devices, BC/DR
Archival, Long-term Storage	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, unstructured files, scientific, video, movies, audio, collaboration, social media history, archive cloud applications, video surveillance and security system history and archives, off-site media storage, remote data vaults, BC/DR

Artificial Intelligence (AI) and Metadata to Improve Data Classification

If data is the new currency, then storage is the new bank. IT staffs are under increasing strain as the volume and complexity of managing daily workloads defy the traditional approach of simply adding more drives 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 and mirror 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. Who access it? How often is it accessed? What kind of file is it? What type of data 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. The storage industry is anxious for AI tools, but they won't suddenly appear - rather they will evolve and improve their capabilities over time.

Acquisition Cost and TCO Comparisons for a Tiered Storage Implementation

The example below compares the initial acquisition cost of 10 PB of storage for various tiered storage scenarios. The ASPs (average selling price) are calculated from 2020 projections for each tier and serve only as an example as prices can vary widely. Businesses should input their own respective storage cost metrics or use actual price quotes from your storage vendor(s) to determine costs for a specific tier. The two disk tiers are combined and averaged for simplicity. The percentage of data in each tier uses the industry averages from the Data Classification Model described earlier.

Total Capacity	**Ex: Acquisition Costs for 10 PB Capacity**			
Storage Tiers	**Tier 0 (SSD)**	**Tiers 1&2 (disk)**	**Tier 3 (tape)**	**Total**
Optimal data allocation % by tier	10% (1PB)	30% (3 PB)	60% (6 PB)	10 PB
ASP/TB by tier (2020 est. price)	$140.00	$40.00/$30.00 (assume $35.00)	$4.00	NA
Total 10 PB price per tier if all data on each tier	$1,400,000	$350,000	$40,000	NA
Total optimal price (SSD and disk tiers only)	10% $140,000	90% $315,000	0%	100% $455,000
Total optimal price (SSD, disk and tape)	10% $140,000	30% $105,000	60% $24,000	100% $269,000 (-41%)

The more data that can move to tape (tier 3), the greater the cost savings over an all disk approach. If 10 PB was stored in tier 1, the total cost would be $1,400,000. If 10 PB was stored completely on the two disk tiers, the cost would be $350,000. By including the tape tier to optimize costs, the cost would be $269,000 reducing the cost by $186,000 or 41%. Moving as much data as possible to the tape tier (3) pays the largest economic dividends. The real cost (TCO) benefits of tiered storage increases over subsequent years as the operating expenses all quickly dwarf the upfront costs to purchase and deploy. Any way you look at it, the larger the storage pool, the greater the tiered storage benefits become. For TCO comparison, the chart below compares automated tape, low-cost disk, and the cloud for tier 3 storage.

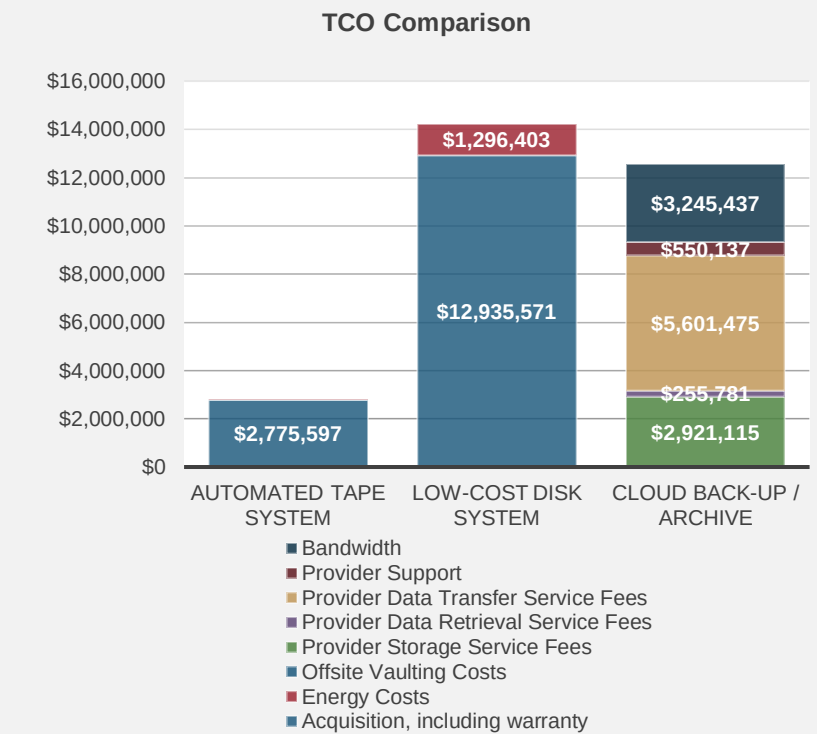

TCO Key Assumptions **20 PB of storage**

- 30% CAGR
- 12% of data retrieved year
- LTO 8 Cartridge Price of $112.25
- Technology refreshed in 5 years
- Energy cost of $.105 based on commercial cost per US Energy Information Agency
- TCO includes acquisition, energy, maintenance, cloud storage, network and technology refresh costs
- Tape TCO includes cost of fully automated (robotic) tape library and drive systems

10 Year TCO Results

- Tape Savings versus Disk Storage – 80%
- Tape Savings versus Cloud Storage – 78%

The Data Avalanche is Underway

Even as CIOs struggle with the exploding growth of disk farms, many continue to maintain expensive tier 1 or tier 2 disks often half full of data which has little or no activity for years. Obviously, as data center storage farms grow, few can afford to sustain this degree of inefficiency and the value of tiered storage is evident. The unknown magnitude of the Internet of Things (IoT), the big data pile up, mobility, surveillance, social media, and countless new tele-health requirements resulting from the global Covid-19 pandemic are all key reasons why data storage volumes are expected to continue to grow at over 20% annually. The IoT, coupled with the arrival of 5G networks which are ~20x faster than 4G, promises to enable countless data generating end points that may number in the billions and trillions.

Conclusion

An effective tiered storage strategy has to balance storage costs, data lifecycle management practices, and storage technology priorities with storage performance. The best choice is to take advantage of a multi-tier storage system that automatically migrates your data to the most cost-efficient (TCO) tiers of storage. Clearly today's storage management challenges will only get larger. The long-standing goal of delivering on the promise "to have the right data - in the right place - at the right time" can finally become a reality. For some enterprises building a tiered storage strategy may still be optional, but that won't be the case for much longer.