

Reinventing Archival Storage



THE STORAGE HIERARCHY




Fred Moore
President
Horison.com

The Zettabyte Era Arrives

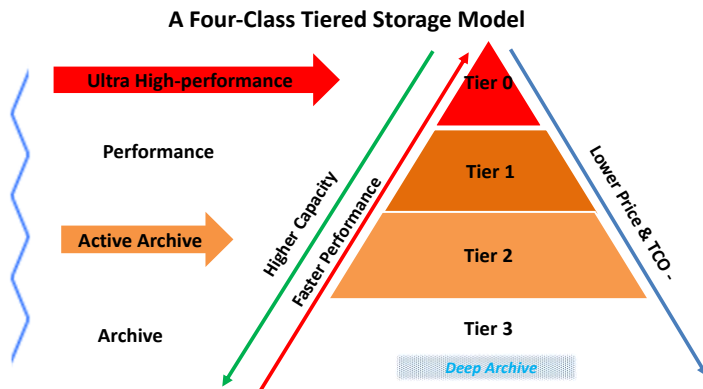
Data creation is on an unprecedented trajectory and data has become the most valuable asset in most companies. For example, in 2006, oil and energy companies dominated the list of the top five most valuable firms in the world, but today the list is dominated by firms based on digital content like Alphabet (Google), Apple, Amazon, Facebook and Microsoft. [IDC's](#) latest report estimates that the Global StorageSphere installed base of storage capacity (not data) is more than doubling over the forecast period and is expected to grow between 2019 – 2024 at a CAGR of 17.8%, resulting in an installed base of storage capacity of 13.2 ZB (1×10^{21}) in 2024, compared with 5.8 ZB in 2019. The amount of data that is actually stored, or utilized storage, is expected to grow at a 20.4% CAGR, from 3.5 ZB in 2019 to 8.9 ZB in 2024.

The storage infrastructure is a major component of most IT budgets. Storage tiering has become the key strategy that enables optimum use of storage resources, reduce costs and make the best use of available storage solutions for each class of data. Balancing a storage architecture to optimize for different workloads while minimizing costs now presents *the* next great storage challenge. In addition, over 60% of all data is archival and it could reach 80% or more by 2024, making archival data by far the largest storage class. Given this trajectory, the traditional storage hierarchy paradigm will need to disrupt itself – quickly.



Balancing Performance and Cost – Tiered Storage Architecture Provides the Greatest Benefit

Today's storage hierarchy is well established and consists of a four-class tiered model. Each tier has its



- The tiered storage model balances price, performance and capacity for the optimal storage solution
- Performance is the key characteristic for Tier 0, price and TCO are increasingly critical for tiers 1,2,3
- Over 60% of all data is archival
- A New *Deep Archive* Tier is Emerging

Source: Horison, Inc.

own unique characteristics and value proposition. Today's storage tiers range from ultra-high capacity, low cost storage at the bottom of the hierarchy to very high levels of performance and functionality and at the top. A multi-tiered storage architecture capitalizes on the significant cost differences of each tier by deploying advanced management tools that move, migrate, and respond to changing lifecycle needs of the data or application. Deploying as much data as possible in the archive tiers yields the most cost-effective

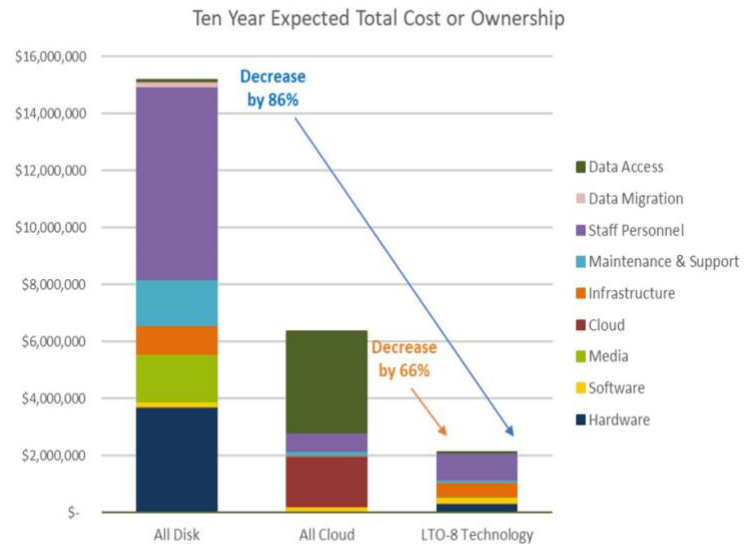
architecture. As storage pools increase in capacity, the business case for implementing tiered storage becomes more compelling potentially saving millions of dollars in storage costs.

A Four-Tiered Storage Model - Key Characteristics and Future Requirements

Storage Tier	Tier 0 Ultra-high Perf.	Tier 1 Performance	Tier 2 Active Archive	Tier 3 Archive, Deep Archive
Key Considerations	Extreme IOPs, access & response time, price not a primary concern, high reliability	High IOPs, online DBs, price balanced with perf., highly sensitive, high reliability	Data not constantly in use, reasonable access times, price is a major consideration	Infinite retention periods, long media life, huge data volumes, seldom accessed, price is most critical
Workloads	OLTP, Res systems, HPC, AI/ML, scientific, real-time simulations, mission critical apps	Mission critical DBs, highly sensitive, Nav systems, ATM, POS, Govt. security, tele-medicine,	Business sensitive apps, backup, active archival, reference data, surveillance, dev and test, ERP, mobile and edge apps, email	Compliance, archives, easily scalable and require minimal remastering cycles, access times of hours or even days can be acceptable.
Time to First Bit	1-10 μ (microseconds)	5 - 10 ms (milliseconds)	Secs, minutes	Secs, minutes, hours...
Throughput (Data Transfer Rate)	550/520 MB/s R/W speed to 3,500/2,300 MB/s R/W speed	160 - 220 MB/s	80 - 220 MB/s	400 MB/s
2020 est. price	\$150/TB	\$30/TB	\$8.0/TB	\$4.00/TB
Price	Not a concern - performance at any cost	A key decision factor	Major concern	Critical concern
Reliability (BER)	1×10^{17}	1×10^{16}	1×10^{16}	1×10^{19}
Media Life	3-5 years	4-5 years	4 - 30 years	50 - 100+ years
Power Usage	2-5 W	6-15 W	6-15 W	Lowest (air-gapped)

For Storage – It Always Comes Down to Price and TCO


While price is and will continue be the single greatest driving force behind storage investments for the foreseeable future, a tiered storage strategy offers the most significant cost benefits compared to any other storage solution for nearly every workload. The attributes of long shelf-life, data integrity and extremely low acquisition costs play heavily in the TCO (Total Cost of Ownership) as data retention periods can exceed 100 years encouraging archival data to reside in the lowest cost storage tier. The [adjacent TCO study](#) estimates total storage costs over the next decade showing disk systems costing as much as 7 times that of a tape system and a cloud system costing more than 3 times that of tape. Getting the right data in the right place has great benefits.



Data Center Cost Models Reveal the *Compelling Value of Tiered Storage*

Industries optimize for different use cases as all workloads are not created equal. The diagram below applies costs of each storage tier to model five different data center workloads to store one exabyte (1x10¹⁸ bytes) of data. Moving data to the optimal tier for each unique workload offers tremendous cost savings. Data centers have to balance unique, highly variable workflows with legacy and current applications while adopting new paradigms that exploit a variety of storage devices, systems, and architectures. Some are computing intensive; others are data and I/O intensive. Many collect and store large volumes of potentially important data, but it may remain inactive for years waiting to be mined. In other cases, raw data is collected and immediately processed, then stored for years just in case there’s a need for further processing or deeper analysis. In any case, optimizing the storage tiers is the key.

Five common data center workload models and Hyperscale Data Centers (HSDCs) are shown using the pricing assumptions below for allocating 1 exabyte of storage capacity across the four tiers.




UHP = Ultra High-Performance Tier
 P = Performance Tier
 AA = Active Archive Tier
 DA = Archive, Deep Archive
 TD = Total Data Stored (Terabytes)

**Ex: 1 exabyte capacity
 Optimizing Data Pyramid
 Budget**
 $(UHP\% * TD * UHP \$/TB)$
 $+ (P\% * TD * P \$/TB)$
 $+ (AA\% * TD * AA \$/TB)$
 $+ (DA\% * TD * DA \$/TB)$

High-performance workload (Heavy Compute):


These HPC workloads have a high-performance computing intensive architecture, where thousands of servers are networked together into a cluster. Software programs and algorithms are run simultaneously on the servers in the cluster with high performance storage systems to capture the output. HPC supports compute intensive applications requiring fast random access, block storage for pattern recognition that simulate future consequences and predict outcomes. HPC organizations use extensive analytic services (OLAP) to understand the vast amounts of data across complex hybrid environments.



Class	Data Distribution	Cost
UHP	70%	\$105,000,000
Perf	10%	\$3,000,000
AA	10%	\$600,000
DA	10%	\$100,000
Total	100%	\$108,700,000

Transactional workload (Heavy reads, writes and updates):


These workloads are response time critical business processes such as OLTP, real time communications, reservation systems, order processing optimized for heavy reads, writes and updates on small records, and ability to modify data in place. The financial industry uses heavy transaction-based services for online banking transactions, ATM, POS, credit card charges, audit and communication logs.



Class	Data Distribution	Cost
UHP	30%	\$45,000,000
Perf	20%	\$6,000,000
AA	20%	\$1,200,000
DA	30%	\$300,000
Total	100%	\$52,500,000

Traditional Data Center workload (Heavy Database):


This is the most common type of data distribution model. The de-facto standard database workload must be tuned and managed to support the service that's using that data. Health Care and Life Sciences use databases for electronic medical records, images (X-Ray, MRI, or CT), genome sequences, records of pharmaceutical development and approval fit this profile.



Class	Data Distribution	Cost
UHP	10%	\$15,000,000
Perf	10%	\$3,000,000
AA	20%	\$1,200,000
DA	60%	\$600,000
Total	100%	\$19,800,000

Batch workload (Business Support Activities):


These workloads are common and are designed to operate in the background and tend to process larger and more sequential volumes of data. Generally, batch workloads are seldom time sensitive and are executed on a regular schedule. Examples include daily reports from ERP, inventory, email, payroll, monthly or quarterly billing or compiling the results of months of online reports. Insurance processes many long-term accidents records and images, health claims, disputes, and payment history in batch mode. These workloads may require considerable compute and storage resources when in use.



Class	Data Distribution	Cost
UHP	5%	\$7,500,000
Perf	15%	\$4,500,000
AA	30%	\$1,800,000
DA	50%	\$500,000
Total	100%	\$14,300,000

Data Intensive workload (Heavy Archive):

Data-intensive applications devote most of their processing time to I/O and movement and manipulation of data as search times and data transfer speeds, not processor speeds, limits their performance. For example, the MLB Network archives over 1.2 million hours of content, which is indexed and stored with infinite retention periods. The M&E industry relies heavily on digital archives as they provide raw production footage. For movie production, it is common to have workflows that need access to archived digital assets at some point in time after the movie is made. Most M&E content is never deleted. Surveillance, medicine, finance, government, physical security and social media also fit this model.



Class	Data Distribution	Cost
UHP	5%	\$7,500,000
Perf	10%	\$3,000,000
AA	10%	\$600,000
DA	75%	\$750,000
Total	100%	\$11,850,000

Hyperscale Data Center workloads (Compute, Storage, Backup, and Archive cloud services)

With nearly [600 data centers](#) worldwide and 40% in the US, the HSDC represents the fastest growing data center segment today. Most are large CSPs (cloud-service providers) offering highly distributed compute, storage, archive and backup services and can profile any of the five models above. They constantly re-engineer storage strategies to manage extreme data growth, enormous energy consumption, high carbon emissions, and long-term data retention requirements. HSDCs are struggling with insurmountable growth of disk farms which are devouring IT budgets and overcrowding data centers forcing data migration to lower cost archival solutions. Obviously, few data centers can afford to sustain this degree of inefficiency. HSDCs challenges are extreme at scale as they are accelerating the need for an archival breakthrough.

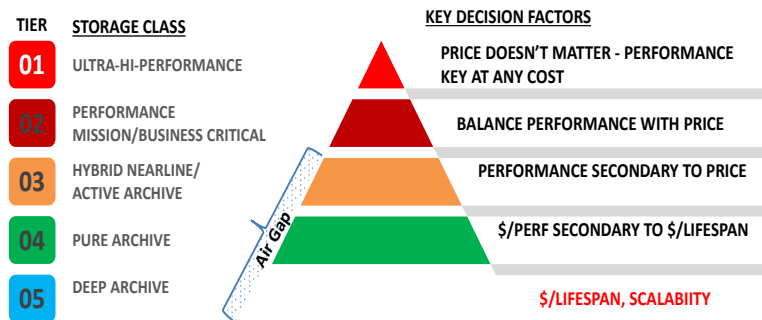
Summary - Tiered Workload Models	Cost for 1 Exabyte	Example Industries
High-performance	\$108,700,000	HPC, DOE, DOD and Government Data Centers, Scientific Research, Weather Models, GIS, Genome Mapping, ML, AI
Transactional	\$52,500,000	OLTP, Online Banking, ATM, POS, Credit Card Processing
Traditional	\$19,800,000	Manufacturing, Financial Services, Insurance, Health Sciences, Online Advertising, Clickstreams, B2B, Direct Sales
Batch	\$14,300,000	ERP, Daily Summaries, Monthly Billing, Periodic Reports
Data Intensive	\$11,850,000	M&E, TV Programming, Sports Archives, IoT, Mobile Apps, Social Media, Surveillance, Search Engines
Hyperscale Data Centers (HSDC)	NA - varies by services	Amazon, Apple, Baidu, Digital Realty Trust, Facebook, Google, Huawei, IBM, Microsoft, Tencent

The Future Looms Huge - How Do We Get from Today to Tomorrow?

Today's storage hierarchy will need to evolve providing more cost-effective solutions to address the archival avalanche. The emerging *Deep Archive* (cold) storage tier will need a new breakthrough solution to provide highly secure, long-term storage and require minimal remastering cycles. Remastering is a costly, labor-intensive process that typically migrates previously archived 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, long media life expectancy, minimal environmental footprint, and a pricing approaching \$1/TB, a price significantly below all other competitor's offerings. Fast access times are not critical for the Deep Archive tier

THE FUTURE OF THE STORAGE HIERARCHY

By 2025 the Tiered Storage Hierarchy Will Evolve



- A new *Deep Archive* storage paradigm will be required.
- Longest life media (100+ years), minimal remastering effort and environmental footprint.
- Random and sequential access needed for mass storage.
- Purchase price ~\$1.0 \$/TB, negligible TCO.

BY: HORISON, INC.

What's Ahead Current Storage Technologies?

For **solid state** (semi-conductor) storage, the roadmap shows significant progress through 2029. The increase in 3D NAND flash layers is leading to huge capacity increases at the die level. Leading suppliers are preparing for 128-layer 3D NAND, while outlining a roadmap to a staggering 800+ layers by 2030. Technology and cost challenges exist but are not expected to be insurmountable for the planning period. Solid state storage has broad market appeal and is heavily used from personal appliances to HSDCs.

For **magnetic storage** (disk and tape), the [INSIC roadmap](#) indicates challenges exist but are not expected to be too difficult to overcome in basic disk and tape technology through 2029. SSD will continue to take market share and high profit margins from disk at the high end of the storage market, while disk is being squeezed by tape for low cost/TCO archival storage at the low end. Lab demonstrations show HDD capacities approaching 40 TBs and native tape cartridge capacities reaching 400 TB's are in sight. Single-layer disk areal density growth is slowing while performance gains are negligible. Tape is a sequential access only media and time to first bit can be a concern

For **optical disc**, the price, performance, low capacity and reliability levels, and slow roadmap progress compared to magnetic technologies have kept the Blu-ray optical disc product family confined to the shrinking car and home personal entertainment market. Optical disc has been here for decades but has never been viable in the data center.

Next Up for Storage Innovation?

Multi-layer photonic storage with 16+ layers for terabyte-scale capacities is under development and may offer the greatest promise in the next five years to address the archive avalanche. By delivering random access, extremely long shelf-life, 100-year remastering cycles with a small environmental footprint and low TCO, a truly scalable mass archival storage solution could *finally* emerge. The arrival of a multi-layer recording media at a cost per TB below that of magnetic tape could disrupt the archival paradigm.

Future Speculation – What's Happening in the Labs?

Synthetic DNA based data storage systems have received significant media attention due to the promise of ultrahigh storage density, random access, and long-term stability. DNA claims to offer a storage density of 2.2 petabytes/gram. However, development still must resolve high cost, very long read-write latency and error-rates that render them noncompetitive with modern storage devices. Momentum is building.

Quartz glass media is potentially durable for many decades because of its resistance to heat and weather and is also as hard as porcelain on the Mohs scale. A standard-sized glass disc can store around 360 TBs of data, claiming an estimated lifespan of up to 13.8 billion years even at temperatures of 190°C. That's as old as the Universe, and more than three times the age of the Earth. Progress has been negligible.

Holographic storage has been attempted for years and include nanoparticle-based film that can store information as 3D holograms, improving data density, read and write speeds and stability in harsh conditions. 5D crystal technology is a multi-state optical WORM technology, but the problem is that external UV light can erase the data. That doesn't bode well for the long-term stability of the technology.

Quantum Storage is investigating ways to use quantum physics to store a bit of data attached to the spin of an electron. In future, we could see instant data syncing between two points anywhere, based on quantum entanglement. Quantum computing may ultimately be the truly transformational technology, but "when" remains completely unknown.

Note: *These technologies have been under development for decades and often make headlines but have never made it out of the laboratory into commercial markets. Given the challenges they face, it's unclear when or if any of these will arrive in the foreseeable future. So far there has been more written about these technologies than has been written on them.*

Summary

The storage industry is hungry for new innovative hardware and software solutions that can change the rules of the game. The momentum is building for new Deep Archive storage tier to emerge providing tremendous economic and long-term storage benefits. The successful new tier will effectively address cost, scalability, and security as the most compelling future storage considerations for IT managers. By 2024 archival data will represent between 60 - 80% of all digital data making it the single largest storage class and will approach [\\$11 billion](#) in annual revenues fueling demand for a new storage paradigm that can drive storage costs closer to the \$1.0/TB target level. To meet the archive challenge, disruption may be the greatest differentiator. The key issue now becomes which future technology will be able to successfully reinvent the traditional archival storage tiers, or will businesses be forced to continue using the less than optimal storage solutions of the past?

About the author



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

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