

The Missing Storage Tier



Delivering the Most Sustainable, Cost Effective, Highly Scalable Mass Storage System



Fred Moore, President
Horison Information Strategies
Horison.com

Horison.com
Horison Information Strategies

Preface

Data creation is on an unprecedented growth trajectory as digital data has become the most-valuable asset for most companies. The insatiable AI era is in full swing demanding advanced solutions to effectively contain escalating petascale and exascale storage requirements. Growth projections for data stored worldwide suggest the ~ 25 - 30% CAGR could continue, doubling the amount of data stored approximately every three years. In time AI should markedly improve storage efficiencies keeping growth rates in check or possibly slowing them. This report examines the requirement for a "missing tier" between operational storage and deep archive filling the gap for data that is no longer hot but requires timely (active archive) and unpredictable (dynamic archive) access capabilities. It explores the business and technical drivers behind this requirement, evaluates existing storage technologies, and outlines the characteristics needed for a next-generation secondary storage platform as storage requirements are fundamentally changing.

Satisfying the *missing tier* has become the *new frontier* for data storage.

Table of Contents

Executive Summary	1
1. The Strategic Challenge	4
1.1 Data Growth and the AI Revolution	
1.2 Data as a Long-Term Strategic Asset	
1.3 Why Existing Storage Architectures Are Reaching Their Limits	
2. Active Archives Currently Address the Missing Tier	4
2.1 Evolution of the Storage Hierarchy	
2.2 Active Archive Architectures	
2.3 Benefits and Limitations of Current Approaches	
3. Traditional and AI Data Lifecycle Access Patterns Are Different	5
3.1 Traditional Data Lifecycle Models	
3.2 AI-Driven Data Access Patterns	
3.3 The Emergence of the Dynamic Archive	
4. Understanding Lifecycle and Access Pattern Charts	7
4.1 Data Value Over Time	
4.2 Access Frequency and Predictability	
4.3 The Growing Importance of Historical Data	
5. Missing Tier Architecture Requirements	8
5.1 Three secondary storage tiers today	
5.2 WORM (active archive), WORSE (archive), and WORN (deep archive) lifecycle pattern	
6. Where Does the Missing Tier Fit?	9
6.1 The Storage Hierarchy Defined	
6.2 Characteristics of the Missing Tier	
6.3 Operational and Economic Requirement	
7. Designing the Missing Tier: What Will Be Required	10
7.1 Scalability	
7.2 Cost Efficiency	
7.3 Accessibility and Availability	
7.4 Durability and Longevity	
7.5 Sustainability and Energy Efficiency	
7.6 Security and Data Protection	
7.7 AI Readiness	

8. Today’s Storage Technology Challenges From Here to Where?... 11

- 8.1 Flash Storage Technologies
- 8.2 Hard Disk Drives (HDDs)
- 8.3 Magnetic Tape Systems
- 8.4 Optical Storage and Long-Life Media
- 8.5 Cloud Object Storage
- 8.6 Comparative Analysis of Current Technologies

9. Mapping the Missing Tier to Existing Technologies 13

- 9.1 Requirement-to-Technology Analysis
- 9.2 Technology Capability Matrix
- 9.3 Identifying the Gaps

10. The Path Forward 14

- 10.1 Emerging Storage Technologies
- 10.3 Future Active Archive Architectures
- 10.4 Recommendations for Industry and Enterprises



Charts and tables



Chart 01 p6 Data Lifecycle and Access Patterns (*Before and After AI*)

Chart 02 p9 Designing The Missing Tier – *What Will Be Required?*

Chart 03 p10 Secondary (Archival) Storage Model – 2026 and Beyond

Chart 04 p11 The Missing Storage Tier

Chart 05 p13 Today’s Technology Challenges - From Here to Where...?

Chart 06 p14 Mapping the Missing Tier to Existing Storage Devices

Appendix p15

Chart Vertical Market Failure

Table Potential Storage Developments

Table Readiness Requirements for Current Technologies

Chart The World of Digital Data

Chart Active and Dynamic Archives

Table Archive Tiers in Secondary Storage

Executive Summary

The data storage industry is at a turning point. Unabated AI growth and rising data volumes are exposing the limits of today's storage infrastructure and creating demand for a new secondary storage tier. Enterprises and hyperscale operators need a lower-cost middle layer between expensive primary storage and slower, rigid deep archive for secondary data that still has business, compliance, analytical, operational, or AI value. At the same time, a Vertical Market Failure (VMF) in magnetic storage is well underway as HDD and tape developments slow, and roadmaps fall short of future capacity, cost, and access time needs. Supplier options have also narrowed sharply, from 120 HDD makers to 3, from 40 tape drive makers to IBM alone, and 20 tape media cartridge makers to just 2. The result is a widening storage gap, and an urgent need for a new storage architecture. For the first time in nearly two decades the prices per TB for SSD, HDD and tape have risen as the once predictable and steady price erosion model has reversed its course. As a result, enterprise storage architectures must prepare to enter a new and demanding era.

Without data, AI is useless. AI adoption has accelerated data creation, retention and access demands, and long-term digital preservation is forcing organizations to rethink the way they store and preserve information that is no longer hot (high activity) but remains too valuable to bury in deep archives. The VMF dynamic continues to stifle any breakthrough innovation needed to address the arrival of extreme storage demand at scale. Instead of using this time to address VMF, both hyperscalers and storage manufacturers are focused on fulfilling short-term demand, near-term financial optimization, debt reduction and capital return strategies that, on their current trajectory, leaving long-cycle for future innovation unfunded.

The missing data storage tier has been a recognized architectural gap for years, but the rapid rise of AI has made it impossible to overlook. Enterprises now need a secondary storage layer for data that is no longer continuously active yet must remain searchable, reusable, accessible, and economically practical. This missing tier bridges performance-oriented primary storage, where access is measured in milliseconds, and retention-oriented archival storage, where retrieval may take minutes, hours, or even days. Its role is to deliver access measured in seconds, including random-access capability once the media is available, at a materially lower cost than premium storage and with far more predictable retrieval than traditional deep archive.

The storage industry lacks a committed next-generation storage technology direction. If this issue continues to be unaddressed, both data storage suppliers and hyperscale buyers will lose. If there are no alternatives available at much lower total cost of ownership (TCO) numbers than current offerings, then the suppliers will remain stuck with only derivative new product introductions with no incentive to invest in novel, long lasting technologies. As a result, the *current* storage technology hierarchy has a widening gap between expensive production storage and slower, less flexible archival systems. This gap or "hole in the hierarchy" describes **the missing data storage tier**.

The Strategic Challenge

Historically, organizations optimized storage primarily for cost reduction. Today, they must also optimize for future value extraction. Data that may appear dormant today could become critical for AI model training, regulatory analysis, [digital twin](#) development, or monetization opportunities tomorrow. The challenge is that existing storage tiers force organizations to choose between performance and cost. There is growing demand for a storage layer that delivers near-online accessibility at economics much closer to archive storage - the rise of the missing tier.

Current storage architectures often force an inefficient tradeoff: keep this data on costly performance storage or move it into an archive that can be too slow, too rigid, or too difficult to search when the data is needed again. AI intensifies this problem because it depends on large volumes of accessible, well-governed historical data. Designing a storage architecture to effectively manage all types of data through increasingly longer and more dynamic lifecycles for a wider variety of workloads while minimizing costs presents the next great generational storage challenge. Satisfying this challenge can finally fill the missing tier in enterprise storage architecture.

Active Archives Currently Support the Missing Tier

Archives are now active data sources, not just data dumps. Fueled by AI, a paradigm shift from cold data to an active archive is underway and exposes the missing tier. The active archive concept has emerged to address two long-standing storage technology limitations: tape is sequential and has longer recall and access times normally in minutes, while Nearline HDDs have millisecond access time, but HDD repositories are more costly, energy-intensive, and non-removable.

Active archives combine SSDs or nearline HDDs serving as a cache buffer with tape libraries, private cloud and or public cloud containing archival data behind an intelligent data management layer. This software layer presents the different technologies as one virtual storage pool and applies policies across the data lifecycle. The main goals are to free expensive primary capacity, reduce storage cost, lower energy use, and move cool or cold data to lower-cost, more energy-efficient platforms.

Active archive systems store and retrieve data based on lifecycle stage, access frequency, metadata, user-defined policies, and workload demand. Newly reactivated or frequently used content can migrate from secondary storage to SSD or HDD for faster random access while lower-activity or inactive content can remain stored on tape or cloud archive indefinitely while minimizing energy consumption. The active archive improves practical access to archived data, but it adds operational overhead when large datasets must be moved, copied, recalled, migrated or re-staged across multiple storage tiers. For this reason, the ideal missing-tier architecture should reduce unnecessary data movement between tiers and support active and dynamic archive, bulk archive, and deep archive requirements within a more closely integrated storage subsystem.

Traditional and AI Data Lifecycle Access Patterns are Different

Chart 01 compares the traditional data lifecycle and access pattern models before AI and after the arrival of AI. In the traditional data lifecycle model, data begins as active online data, transitions into an active archive, and eventually becomes long-term, cold, and persistent archival data. The long-tail probability of access $P(a)$ model often guides storage architecture design, retention planning, access expectations, and tiering policies for environments that manage large volumes of archival data without AI.

AI pipelines often ingest large volumes of raw data before its future value is known. Over time, these datasets may be reused for training, retraining, evaluation, bias testing, compliance, new model development, or retrieval-augmented generation (RAG), which improves responses by searching relevant external data. As a result, the missing tier must provide predictable retrieval, support sudden bursts of higher activity, maintain searchable metadata, enable policy-based data movement, and reconnect easily with primary workflows. Rather than moving in a traditional one-way path from hot to warm, cool, and cold storage, archival data can become active after AI processing and then return to a colder state. This change makes access patterns less predictable and shifts storage planning from the traditional long-tail archive model toward a broader dynamic archive model. Today, HDDs combined with tape in an active archive are needed to deliver a random-access mass storage solution to fill the missing tier; in the future, a single technology could conceivably provide the same functions avoiding data movement between tiers.




Chart 01

Traditional Lifecycle - Historically, the data lifecycle follows a predictable progression:

- Data creation
- Frequent operational access
- Reduced access frequency over time
- Recalls occur but are less frequent
- Long-term archival retention period
- Eventual deletion

AI-Era Lifecycle - AI introduces a fundamentally different secondary storage access pattern:

- Data may become inactive for months or years
- New business opportunities (AI) can be unpredictable can suddenly require large-scale retrieval
- Training datasets need to be rebuilt repeatedly
- Historical information may be combined with newly generated data
- AI data can move dynamically among hot, warm, and cold states – the dynamic archive- rather than progressing through a one-way lifecycle
- The dynamic archive model requires storage systems that can support sub-minute access requests while maintaining low long-term costs

Understanding Lifecycle and Access Pattern Charts

The AI era is demanding advanced solutions to effectively contain zettascale storage requirements. The largest storage category is secondary, archival, and lower-activity data, which presently accounts for roughly 80% of all stored data.

- Roughly 20% of stored data remains permanently active online and is optimally stored on primary storage such as SSD or HDD.
- About 80% of stored data is lower activity and is optimally stored on secondary storage, including tape and nearline HDDs.
- However at least 60% of the low activity data remains stored on primary and more costly (HDD) storage presenting an enormous target market for the missing tier to address.
- An active archive system supports both traditional and AI access patterns.
- AI introduces a more dynamic archive access pattern, where any archival data can become temporarily hot, cold hot again and stay hot for variable time periods.
- Note: Today's magnetic HDDs typically last ~5 -6 years before replacement while magnetic tape drives typically last ~8-10 years or more. Tape media significantly outlasts the tape drive lifetime lasting 30 years or more without degradation.

The probability of access profile, P(a), typically declines as data ages. The traditional model shows an initially high access probability near 100% for newly active data, followed by a steady decline after the first days and weeks. After the 90-to-120-day range, most data have transitioned into archival status, and long-term access probability typically fall into the 0% to 5% range. At this point data can migrate to lower cost secondary storage to free up costly primary storage and reduce energy consumption. AI can reawaken the archives at any time.

- The value of data is no longer measured by its age alone; archived data may regain importance depending on new analysis, context, or changing business needs.
- Data age 0 days - Data is newly created and should remain on active primary storage.
- 1 week to 30 days: P(a) begins to decrease, but data still justifies online availability.
- 90 to 120 days - Most data typically have become cool or archival and can safely relocate on lower-cost storage.
- 1+ years: -Data has reached archival status, with access expected to be infrequent.
- 10+ years - Data resides in deep archive with rare access requirements.
- 100+ years - Retention planning requires persistent, immutable media capable of supporting century-scale preservation time capsule storage to avoid frequent remastering.

Designing the Missing Tier – What Will be Required

The missing tier should provide a unified secondary storage layer that supports active archive, standard archive, and deep archive workflows without requiring excessive tier to tier data migration between independent storage systems. The missing tier should also enable a metadata-rich, policy-managed object storage that can operate across cloud-native, on-premises, and hybrid environments in each type of data center (see [table](#) below).


Data Center Type	Typical Scale	Primary Purpose	Key Characteristic
Hyperscale	100+ MW	Massive AI workloads, major cloud providers, HPC	Massive scalability
Enterprise	1-10 MW	Internal operations	Owned and operated by single organization
Colocation	10-100 MW	Shared hosting for multiple clients	Customers lease space, power, cooling
Edge	1-10 MW	Low-latency local processing	Located near end-users or data sources

Key requirements include the ability to deliver predictable retrieval behavior for dynamically reactivated data, random access, improved sustainability by reducing dependence on premium storage and minimizing idle energy consumption. Random access storage is important because it lets a system retrieve specific data directly in milliseconds instead of reading through everything sequentially that can take minutes or hours. Chart below identifies the critical design requirements for the missing tier.

Designing The Missing Tier – *What Will Be Required?*

Meeting Soaring Secondary Storage Demands and Beyond

The Missing Tier - Critical Architecture Design Requirements



Source: Horizon Information Strategies, LLC

Chart 02

Critical design points to fill the missing tier include six key dimensions:

- price (lowest \$/TB and TCO),
- performance (random access, eliminate inter-tier data movement, hot and cold workloads),
- capacity (seamless scaling, rack-scale, legacy emulation),
- reliability (BER $\geq 10^{20}$, media resilient to temperature, humidity, radiation, EMP, saltwater)
- security (air gap, encryption, WORM, removable media, physical vault),
- sustainability: (Lowest CO₂e footprint, power consumption, water use, and e-waste).

The chart below presents the 2026-and-beyond secondary archival storage model. In today's active archive environment, most enterprise data move out of primary storage and into a dynamic, multi-tier secondary archival stack. Industry analyst estimates indicate approximately 6.4 ZB of installed storage capacity by year-end 2025, with roughly 5.12 ZB, or about 80%, classified as secondary archival storage. HDDs currently account for about 84% of the world's installed digital storage capacity, SSDs about 12%, and tape roughly 4%. Unfortunately, an estimated 60% of lower-activity archival data remains in the wrong place—on costly, energy-intensive HDDs. Moving low-activity data from expensive primary or nearline storage to an optimized missing-tier solution may represent the largest TCO reduction opportunity in enterprise storage. The chart below separates secondary storage into three tiers based on where data should be optimally stored according to its activity profile. All three tiers - active archive, bulk archive, and deep archive workflows to minimize or avoiding data movement should ideally be addressed in designing the missing tier.

The Missing Storage Tier


Secondary (Archival) Storage Model – 2026 and Beyond

Three Tiers Currently Address the Secondary Storage Stack

- Secondary Storage Currently Has Three (3) Distinct Tiers
- Data Moves Between Tiers (overhead)
- Can the Missing Tier Do it All?

Archival Data is Mostly Write Once

- WORM – Write **O**nce, Read **M**any
- WORSe – Write **O**nce, Read **S**eldom
- WORN – Write **O**nce, Read **N**ever



Source: Horizon Information Strategies, LLC.

Chart 03

- Active Archive (~15%): Nearline HDD or SSD for data that is written once but read many times, also described as WORM — Write Once, Read Many.
- Archive (~60%): Robotic tape libraries, nearline HDD, and other media for data that is written once and read seldom, described as WORSe — Write Once, Read Seldom.
- Deep Archive / Time Capsule (<5%): TBD advanced long-retention technologies such as tape, ceramics, photonics, 3D media, glass, holographic storage, DNA, and future media TBD for data that is written once and rarely or never read, described as WORN — Write Once, Read Never.

Where Does the Missing Tier Fit?

The missing tier fits between hot operational storage with access times in milliseconds and secondary, archival storage with access times typically ranging from minutes to hours. It is a secondary storage layer for data whose immediate access value has declined but whose future business, analytical, legal, compliance, or AI value remains untapped but may have significant potential value. The missing tier should support searchable metadata, predictable retrieval, durable retention, and minimize movement without requiring organizations to overuse premium storage or accept the limitations of deep archive.

The Missing Storage Tier




Chart 04

In practical terms, the missing tier bridges performance-oriented primary storage (measured in milliseconds) and retention-oriented storage (measured in minutes, hours or days) providing access time in seconds for random-access capabilities once the media becomes accessible. It gives enterprises a place to hold cool and cold data that may need to be suddenly reprocessed, restored, analyzed, or governed over longer periods. It also reduces unnecessary movement between tiers by allowing data to remain accessible in a cost-optimized state until policy, workload behavior, or business need justifies promotion or demotion.

Satisfying the next-generation missing tier architecture will require several critical requirements:

- **Massive Scalability**
 - Support for petabyte, exabyte, and eventually zettabyte-scale environments.
- **Low Cost**
 - Economics approaching archival storage rather than primary storage.
- **Fast Availability**
 - Data should remain near-online (seconds access time) without lengthy restoration processes.
- **Long Media Life**
 - Storage media should support retention measured in decades rather than years.
- **High Durability**
 - Protection against corruption, hardware failure, natural disasters, and cyberattacks.
- **Sustainability**

- Significantly reduced energy consumption and environmental impact.
- **AI Readiness**
 - Support for large-scale retrieval operations and recurring access patterns.
- **Removability**
 - Eliminate energy consumption when media is not in use,
 - For disaster recovery purposes - enables media to be safely and more quickly transported in case of a fire, flood, hurricane, electrical outage, natural disaster or acts of violence,
 - Provides an air-gap for additional layer of cybercrime protection.

Today's Storage Technology Challenges – Can They Fill the Missing Tier?

The missing tier will require removable, random-access mass storage with compelling energy efficient economics for secondary storage systems. It should also support permanent preservation use cases where durability, cost, governance, and accessibility must coexist rather than compete. The key question is whether one technology can meet all secondary storage needs, or whether the future will continue to require a coordinated stack of today's current storage technologies HDD, SSD, tape, or optical, with or without an emerging, novel long-duration media. The chart below compares existing storage technologies and shows that each roadmap faces a different constraint: NVM is limited mainly by price, HDD by access density, power, TCO, and roadmap funding, tape by initial access time and lack of random access, and optical storage by uncertain price, performance, capacity, reliability, and throughput. Overall, it reinforces the argument that no current technology alone can fully solve the future archival/storage challenge, creating the need for a new or hybrid approach to fill the missing storage tier.

Today's Technology Challenges - From Here to Where...?





Current Technology 	Recording Technique	Roadmap Capability	Future Developments	Challenges
NVM (SRAM, DRAM, MRAM, NAND Flash, PCM)	Electronic Charge 	Aggressive development, multiple technologies emerging, CXL, NVMe	Multi-layer 3D stacking (500+), HBM, HBF, silicon photonics, non-volatile?	Price, supply
HDD	Magnetic Field 	Performance and capacity roadmaps falling behind	HAMR, SMR?, 10 TB/platter?, add more-platters, multiple actuators, bit patterned media (BPM)?	Access density (IOPs) limited, TCO, high power consumption, \$/TB/watt, CO ₂ roadmap funding?
Tape	Magnetic Field 	Well defined and sustainable capacity growth, 900+ TB demo, high patent activity	Advanced media, Strontium Ferrite (SrFe), Epsilon Ferrite (ε-Fe ₂ O ₃), improved search capability, 64 tracks?	Initial access time, no random access, No consumer market
Optical storage	Light Beams 	Minimal progress to date compared to magnetics, future development activity increasing	Femto-second lasers, ceramics, glass, photonics, holographic, long-life media, EMP proof	Price, performance, capacity, reliability, throughput TBD? Source: Horizon Information Strategies, LLC

Chart 05

A diverse group of data storage technologies remain under development and are aiming at filling the missing tier between. The near future is unlikely to witness one single replacement for the missing tier. The biggest opportunity will include a new low-energy, random-access, removable, highly scalable mass storage technology with a compelling TCO.

Future developments may include a new tape format, ceramics, photonic storage, glass, holographic storage, DNA storage, and other ultra-durable media. These are aimed at ultra-high capacity, very long retention, low idle energy, and preservation use cases, but they are still in the laboratory with commercial availability unknown. Some of these developments may fit better in the deep archive / time capsule layer until write speeds, reader cost, automation, and ecosystem integration significantly improve. Nonetheless this technology gap explains growing industry interest in new storage approaches designed specifically for long-term, AI-accessible data preservation. Promising directions include:

- Advanced optical storage technologies
- Long-life immutable media
- Intelligent data placement based on AI-driven analytics
- Hybrid active archive architectures
- Ultra-dense cold storage systems
- Automated retrieval and orchestration platforms

Mapping the Missing Tier to Existing Technologies

Chart 06 below maps the missing tier storage requirement to *existing storage devices* and shows that no current technology by itself fully satisfies all missing tier requirements. SSDs and HDDs are needed for performance-oriented needs such as fast access time, high data rate, random and sequential access, and support for hot and cold workloads, while tape is best suited for lowest price, lowest total cost of ownership, low energy use, long-life reliability, air-gap security, removable media, and immutability. Power, cooling, space, water use, and e-waste are becoming core design goals for any secondary storage development given the sheer volume of data addressed. Future architectures will favor media and systems that reduce or eliminate idle energy consumption, avoid unnecessary data movement, and place data on the lowest-energy tier that still meets access requirements.

Today's storage stack requires a mix of SSD/HDD and tape with considerable data movement between tiers to meet these needs. The optimal missing secondary storage tier would combine these benefits into a single scalable storage layer that supports low cost, low energy, strong security, performance, long media life, removability, immutability, and a sustainable development roadmap. Chart 06 suggests that HDD and tape technologies are being forced to compensate for growing storage demand and lack of a single solution to address missing tier requirements.

Mapping the Missing Tier to Existing Storage Devices

Missing Tier Requirement	SSD/HDD	Tape	Missing Tier
Lowest Price and TCO		✓	✓
Lowest energy		✓	✓
Access time	✓		✓
Data rate	✓		✓
Random and sequential access	✓		✓
Mass storage, highly scalable			✓
Hot and cold workloads	✓		✓
Minimize data movement			✓
Reliability, long-life media		✓	✓
Air gap, security features (for DR)		✓	✓
Removeable media		✓	✓
Sustainable roadmap	VMF	VMF	✓
Immutability		✓	✓

- Multiple storage devices (SSD,SSD, tape) currently required to meet missing tier requirements requiring data movement between tiers
- HDD and/or SSD needed to meet performance and random-access requirements
- Tape libraries needed to provide removability, air gap, immutability, greatest energy efficiency
- VMF forcing HDD and tape to play catch-up to satisfy storage demand

Source: Horison Information Strategies, LLC.

Chart 06

The ideal missing tier should consist of a removable random-access media, and most likely a robotic library, that uses no energy when idle, unlike HDDs or SSDs. Ideally the media wouldn't have to rotate, spin or move to be accessed further saving energy, rather the R/W mechanism could cover the entire surface of the media with minimal energy consumption. The company that satisfies the most missing tier requirements will find the door to addressing over 80% of the worlds digital data opening for them.

The Path Forward

AI has broken the assumption that archived data remains cold by turning secondary and archival storage into a dynamic, unpredictably accessed tier. Yet no current technology delivers seconds-level access at archival economics with zero idle energy. Closing this gap is now the central long-range challenge for storage architecture. Filling the missing tier will require deliberate architectural redesign and, quite possibly, novel media. Enterprises must reduce reliance on premium capacity for aging data, lower energy consumption, deliver predictable low-cost access, and improve readiness for AI reuse, compliance, analytics, and operational recovery. The missing tier is not simply an archive replacement; it is an accessible preservation layer that supports both long-term retention and future reuse. Ideally, one technology would satisfy the full set of requirements needed to complete the tier.


Future breakthrough storage technologies are being developed by companies at various stages of fundraising, but a go-it-alone strategy may not be the only viable path. Nothing in today's market standstill appears likely to release the capital required for long-cycle innovation. Even as

manufacturers reach record revenues and profitability during a multi-year product shortage, much of their capital is being directed toward dividends and stock repurchases rather than next-generation storage development. Under these circumstances, an independent storage industry consortium deserves careful consideration however the path from here to a chartered data storage consortium is unclear. By pooling capital from storage stakeholders, such a consortium could address the Vertical Market Failure (VMF), close the investment gap created in part by recent stock buyback authorizations from public data storage manufacturers, and fund the advanced storage systems needed to build a durable missing tier.

As AI unlocks the value of historical data and accelerates new data creation, demand for a dedicated missing storage tier will only intensify. The need is urgent, the opportunity is enormous, and storage roadmaps are already under pressure. Clearly satisfying the *missing tier* has become the *new frontier* for data storage.

Appendix

Vertical Market Failure



Potential Storage Developments

Several storage technologies have been proposed and are in lengthy development cycles, some are generously funded, but are any poised to make a measurable impact by 2030?

Potential Storage and Missing Tier Developments				
Technology Candidates	Description	Strengths / Value Proposition	Current Limits	Best Fit
Advanced optical storage Folio Photonics, Ewigbyte, SPhotonics	Includes glass, 3D optical storage, photonic media, and next-generation optical lasers.	Very long media life, low or zero power at rest, immutability, removability, random access, and favorable archival economics.	Lower capacity, slower access time, and lower data rates than current magnetic HDD and tape.	Deep archive, time-capsule retention, and potentially active archive if random access and automation improve.
Ceramic storage Cerabyte	Uses durable ceramic-like media for extremely long data retention.	Random access, resistance to heat, humidity, radiation, electromagnetic events, energy efficiency, and environmental degradation.	Commercial maturity and large-scale ecosystem readiness remain developing.	Active archive, archive, deep archive, permanent preservation, sovereign archives, scientific records, cultural preservation, and compliance records.
DNA data storage Twist, Illumina, Microsoft, Western Digital	Encodes digital data into synthetic DNA molecules.	Extraordinary theoretical density and very long retention under proper conditions.	High write/read cost, long latency, limited automation, and very slow technology maturity timeframe.	Ultra-cold, time-capsule archive where data is rarely, if ever, accessed.
Holographic storage Holomem	Stores data volumetrically rather than only on a surface.	High density and parallel read potential.	Commercialization remains difficult; media, drive cost, and reliability remain challenges.	Possible future high-capacity archival or nearline use if economics improve.


Technology	Description	Strengths / Value Proposition	Current Limits	Best Fit
Computational storage Intel, AMD, Samsung, NGD, Scaleflux, Marvell	Moves processing closer to where data resides. Typically, SSDs.	Reduces data movement, latency, and energy consumption.	Primarily an architectural enhancement rather than a bulk media replacement.	AI preprocessing, search, indexing, compression, encryption, and analytics near the storage layer.
Memory-centric and persistent memory Samsung, SK Hynix, Micron, Intel, Kioxia, WD	Includes storage-class memory and non-volatile memory technologies.	Faster than SSDs and useful for hot data, AI workloads, and metadata-intensive systems.	High cost and endurance constraints.	Hot/warm tiers, metadata engines, and acceleration layers—not bulk archive.
Next-generation tape IBM, Fujifilm, Sony	Tape will likely continue as a major archival technology for enterprise and hyperscale.	Low cost per TB, air gap, removability, long media life, low energy at rest, and an achievable long-range roadmap.	Initial access time, sequential access only, and vendor concentration.	Large-scale cool, cold, and deep archive.
AI-managed storage tiers Software – TBD	Uses machine learning to predict data value, access probability, compliance needs, and best optimal placement.	Automates lifecycle movement across SSD, HDD, tape, optical, cloud, and future media.	Remains a future storage management capability rather than a standalone media technology.	Dynamic archive environments where AI can make old data “hot” again.
* Removable random-access mass storage media	A candidate missing-tier concept combining removability with random access and low idle energy.	Access time in secs, removable media, random access, low energy at rest, long life, strong durability, predictable retrieval, reduced tier-to-tier migration, lower cost than HDD.	Requires a commercially viable media, drive, automation, and ecosystem roadmap.	The missing storage tier between active storage and deep archive.



- Removeable, random-access, mass storage may be the most important missing-tier characteristics.
- Ideal characteristics: access time in seconds, removable media, random access, low energy at rest, long life, strong durability, predictable retrieval, eliminate tier to tier migration, and lower cost than HDD.
- Best fit: the “missing data storage tier” between active storage and deep archive.


Readiness Requirements for Current Technologies

Requirement	Flash	HDD	Tape	Cloud Archive
Low Cost	Poor	Moderate	Excellent	Moderate
Fast Access	Excellent	Good	Poor	Moderate
Scalability	Good	Good	Excellent	Excellent
Durability	Good	Good	Excellent	Good
Sustainability	Moderate	Moderate	Very Good	Moderate
AI Readiness	Excellent	Good	Limited	Moderate



Active and Dynamic Archives

Approximately 80% of all Data is Low Activity, Archival or Cold (the Secondary Storage Market)



Archive Tiers	Secondary Storage Evolving Into Three Tiers	Primary Archive Mode
Active Archive <i>Hot</i>	Fast HDD (Active) or SSD (Instant) online random access to dynamic “active archival data”	WORM Write Once Read Many
Archive <i>Cool</i>	Traditional large, lower activity archival storage, big data and data lakes, cool archives	WORSE Write Once Read Seldom if Ever
Deep Archive <i>Cold</i>	Emerging technology tier for permanent, rarely accessed, dark data, often unclassified or untagged, the “golden or master copy”	WORN Write Once Read Never



[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 new growth opportunities for end-users, storage industry providers, and startup ventures. Fred Moore President.

© Horison Information Strategies, Boulder, CO. All rights reserved. [v 3.0]