TRENDFOCUS

Tape and Archive Storage Service

CQ2 '22 Quarterly Update | September 20, 2022

Sustainable Growth of Archive Storage Will Require New Secondary Storage Solutions

Guest contributor, Fred Moore, discusses the potential for vertical market failures, plus additional TRENDFOCUS analysis

Introduction

For the Tape and Archive Storage Service CQ2 '22 Quarterly Update, guest contributor Fred Moore of Horison Information Strategies offers an interesting view on the potential for vertical market failures in tape and HDD storage should the industries not address the investment required to fulfill rapidly expanding archival storage growth. With Mr. Moore's extensive industry experience and knowledge, his analysis backs up his warning to both the tape and HDD industry that investments in the future are required now to support the long-term archival storage needs of the world's data.

The second part of this report provides some historical and forward-looking analysis of tape media prices on a $/GB level across several LTO tape generations. Media costs to hyperscale customers are reported, illustrating once again the power and scale of large cloud customers to consume high volumes of storage capacity at the lowest costs in the market. The timing and ultimate native capacity achieved during the next-generation transition to LTO-10 will determine how quickly the industry can scale capacity cost-effectively.

Finally, with the recent publication of TRENDFOCUS' **Revised Long-Term Forecast** for nearline HDD capacity, the hyperscale archival storage opportunities for spinning disk media have been revised. While economic turmoil and component supply shortages continue to impact near-term cloud storage demand, hyperscale remains the driving force for storage capacity growth across all technologies, challenging industry players on how to fund development to sustainably meet this need. This last point circles back to Moore's article that kicks off this quarterly update.

TRENDFOCUS, Inc.
20660 Stevens Creek Blvd.
Suite 165
Cupertino, CA 95014
T: (408) 996-9903
F: (408) 996-9910
www.trendfocus.com

Mark Geenen
mark@trendfocus.com

John Chen
johnchen@trendfocus.com

Don Jeanette
don@trendfocus.com

Jason Lervold
jason@trendfocus.com

Tape and Archive Storage Service

CQ2 '22 Quarterly Update | September 20, 2022

Table of Contents

The Secondary Storage Market Faces Mounting Challenges – The Gap Widens Between Hyperscale and Enterprise Storage

Contributed by Fred Moore, President, Horison Information Strategies

Fred Moore began a 21-year career with StorageTek as the first systems engineer and concluded as corporate vice president of Strategic Planning and Marketing. In 1998, Fred founded Horison Information Strategies in Boulder, Colorado, a data storage analyst and consulting firm specializing in keynote speaking, executive briefings, webinars, marketing strategy, and business development. Fred served as Editor of Storage for Computer Technology Review Magazine and has written numerous books, articles, research reports and webcasts for the storage industry.

The Hyperscale Data Center (HSDC) secondary storage requirements for archival, cold, and inactive data are soaring, demanding advanced solutions for petascale and exascale archival storage systems. Today just three HDD manufacturers, two tape media manufacturers and one tape drive manufacturer cover nearly all the world's digital storage needs. Three HSDCs, Amazon Web Services, Microsoft Azure and Google Cloud represented 63% of global spending on cloud services in the second quarter of 2022 and grew 42% collectively. In addition, HSDCs manage over 50% of the world's data pushing beyond current secondary storage limits. Neither HDD nor tape is currently positioned by itself to effectively meet the enormous HSDC storage demands that lie ahead. With these challenges mounting, the probability for a Vertical Market Failure (VMF) in secondary storage is increasing.

An Impending Vertical Market Failure?

A vertical market "fails" when transactions within it are too risky, and the contracts designed to overcome these risks are too costly (or impossible) to write and administer. The typical features of a failed vertical market are (1) a small number of buyers and sellers; (2) high asset specificity, durability, and intensity; and (3) frequent transactions". Problems arise when the market has only one buyer and one seller or only a few buyers and a few sellers. These conditions contribute to the widening gap between HSDC storage demands and actual product solutions in the secondary storage market setting the stage for a VMF.

Chart 1
Properties of Vertical Markets

Chart 2
Secondary Storage Pricing Forecast

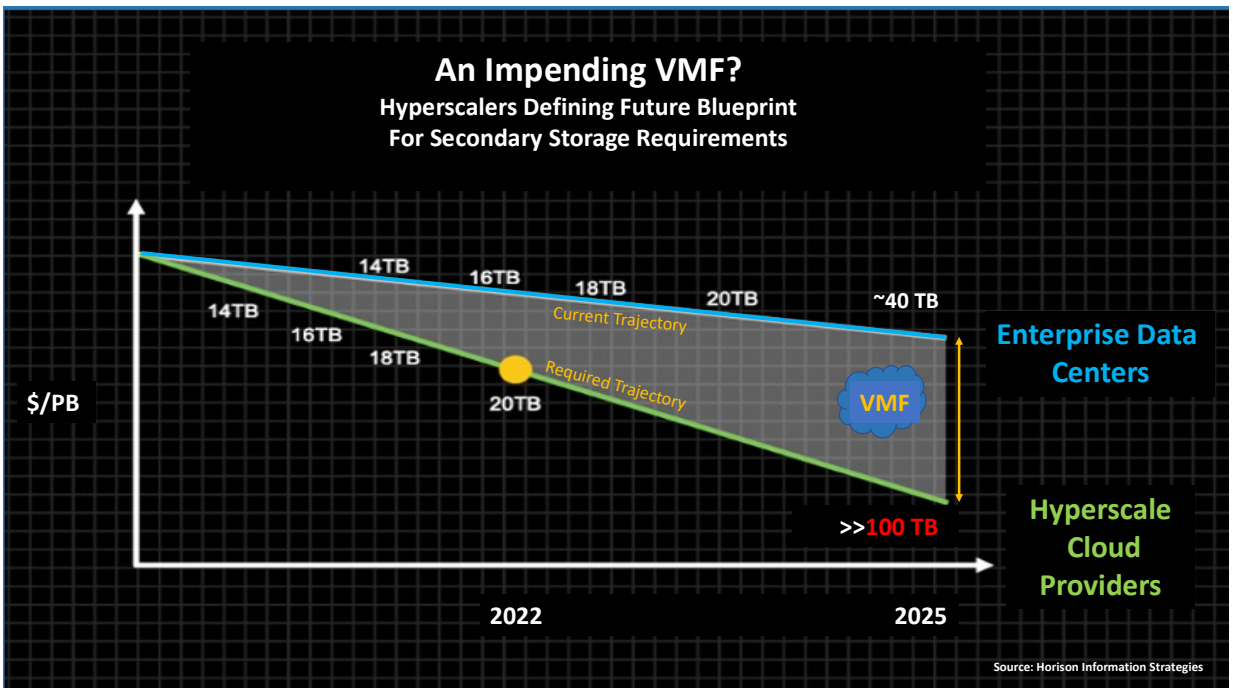

Storage Industry Landscape

Industry estimates can vary but for 2025, market intelligence firm TRENDFOCUS projects the amount of digital data stored to reach ~9.4 ZB. For 2026, the firm projects ~4.7 ZB of storage capacity to be shipped across NAND, HDD, and Tape (native capacity, not compressed). Most of the digital data generated is transient and not permanently stored. Other estimates project that as much as 80% of this stored data (~7.52 ZB) will be classified as archival or cold data defining the magnitude of the secondary storage market, the largest and fastest growing storage segment. Archive data growth rates are projected to range between 25-35% through 2030 though AI, ML, DL and the IoT could present upside. TRENDFOCUS indicates tape capacity shipped growing at 19% annually, but still not effectively capturing the optimal share of the secondary storage demand. Moving inactive data from HDDs to tape, deleting unused, unknown, unidentified, and unwanted data doesn't seem to have much momentum as all types of data continue to pile up. Though tape is currently the most cost-effective archival data solution, most cold data resides on expensive to operate, energy consuming HDDs spinning 7x24x365 adding to corporate sustainability challenges. With over 700 data centers worldwide, HSDCs provide the greatest advantages for efficient and sustainable operations by reducing energy and carbon footprint over the data life cycle.

Shifting device specifications and the slowing rate of HDD and tape technology roadmap development in recent years coupled with HDD and tape storage supplier consolidations have become alarming trends. The HSDC secondary storage needs are growing much faster and are differentiated from traditional enterprise data center archival applications as current "vendor-driven" strategies now lag HSDC demand requirements. HSDCs require

rapid innovation and an individual HSDC may not be able to create the required secondary storage ecosystem. In addition, the development of a new storage device and its format typically takes at least 5 years. This signals that a new portfolio of secondary storage solutions not currently available will be needed to meet future HSDC capacity and performance demands.

The Race to Zero Adds Pressure to Suppliers

Storage negotiations often conclude with price as the deciding factor as storage suppliers strive to reduce prices to gain competitive advantage. Suppliers often count on storage growth to transform the company, but that bet may not necessarily pay off. The storage growth rate must exceed the rate of price erosion to enable profitable revenue growth. With less money to invest, persistent price erosion raises red flags regarding how much erosion storage manufacturers can absorb to remain profitable and invest in the future. Price reduction is critical for HSDCs who leverage their immense bargaining power to drive down supplier prices. However, the race to zero ($0.00/PB) puts supplier margins, future R&D funding levels and innovation at significant risk. Numerous industry consolidations have resulted from the bottom-line pressures from steady price erosion. Therefore, storage manufacturers are forced to scrutinize their cost structures, supply chains and in some cases relax new development projects and schedules.

Sustainability Is a Major Factor

Though HSDCs are very energy efficient, they average of 20-50 megawatts (MW), can often exceed 100 MW, and consume 10-50 kilowatts (KW) per rack as total HSDC power consumption is on another level. Due to the massive energy footprint of HSDCs, climate protection measures have become increasingly important in recent years. HSDCs rely heavily on the internet and research estimates suggest that by 2025, the internet could use as much 20% of all electricity produced and emit up to 5.5% of the world's carbon emissions placing additional pressure on HSDCs to control infrastructure energy consumption.

An Advanced Secondary Storage Model Evolves

If optimally allocated, as much as 80% of all digital data would reside on some type of secondary storage. However most archival data is currently stored in the wrong place and still resides on costly HDDs. For HSDCs this presents a costly, energy hungry and difficult growth strategy paving the way for a new model.

Table 1
Archive Tier Properties

Archive Tier	A New Secondary Storage Model Emerges	Primary Mode
Active	Fast (HDD or SSD) online access to dynamic "active archival data"	WORM
Archive	Traditional large, lower activity, big data and data lakes, cold archives	WORSE
Deep	Emerging technology tier for permanent, rarely accessed, dark data, often unclassified or untagged, the "golden or master copy"	WORN

Tape and Archive Storage Service

CQ2 '22 Quarterly Update | September 20, 2022

Two new secondary storage tiers are evolving to accompany the traditional archive tier as advanced learning applications using AI, ML and Big Data analytics have reawakened the archives. These applications require faster access to analyze archival data but for shorter periods of time as the data will return to archival status after analysis. An Active Archive delivers a big performance boost by integrating two or more storage technologies (SSD, HDD, tape, and cloud) behind a file system providing seamless management for dynamic archival data in a single virtualized storage pool. Expect the Active Archive to become a de-facto standard secondary storage tier over the next few years. The emerging deep archive (Time Capsule) tier will likely be addressed by a new technology and targets data to be stored forever (a golden, immutable copy) but may rarely if ever be accessed. All three tiers are gaining momentum as a new secondary storage model strives to meet increased demands. The chart below describes the flow of data from its creation – to stored tier – to the amount of archival data (~80%) stored in secondary storage.

Chart 3
Secondary Storage Model Attributes in 2025

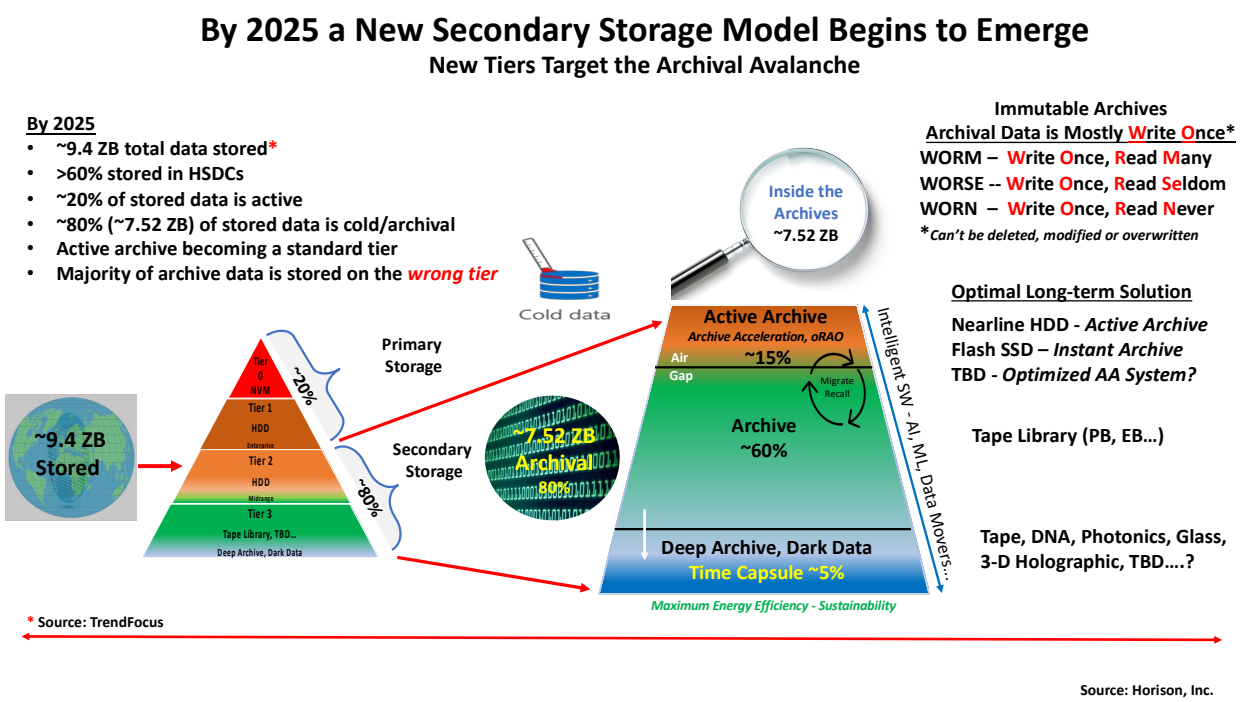

Source: Horison, Inc.

Secondary Storage Scenario (2022)

- The slowing rate of HDD and tape roadmap development in recent years coupled with HDD and tape storage supplier consolidations are setting the stage for a VMF.
- The HDD industry is facing various challenges to increase capacity while delivering any performance gains (IOPs) appear to be minimal.
- Both HDDs and tape are scaling capacity much faster than performance.
- Cost reduction is critical for HSDCs, however the race to zero ($0.00/PB) places supplier margins, future R&D funding levels and storage innovation at significant risk.
- For traditional enterprise data centers, the current rate of HDD and tape product development may be sufficient to meet needs for the near future.
- For HSDCs, data growth and sustainability challenges will eventually become insurmountable without the arrival of a new, cost-effective, scalable, energy efficient secondary storage solution.

Will New and Emerging Technology Developments Help?

Several development activities including photonics, DNA, glass, and holographic are underway attempting to contain the looming secondary storage challenge. Some have been in labs for decades and have significant challenges to overcome, and none appear to be on the short-term horizon for broad-scale customer deployment though photonics and DNA are attracting market awareness. Tape is the incumbent and front-runner in this race and has a complete suite of storage management software in place. The 580TB tape cartridge capacity demonstration from Fujifilm and IBM using Strontium Ferrite magnetic particles, a 29x greater capacity than the current 20TB enterprise cartridge maximum, suggest a long and promising development path lies ahead for tape recording. However, simply scaling capacity without scaling performance will only increase access times, device contention and reduce access density.

Looking ahead, HSDCs will likely use everything the SSD, HDD, and tape industries can manufacture and deliver, along with any other emerging technologies to cost effectively satisfy the demands of the zettabyte era. The ability to support mass production and commercialization of any of these will take time. Fortunately, all these initiatives recognize a major strategy change is needed to get ahead of the soaring storage demand curve as the zettabyte era is well underway. See chart on the following page.

Chart 4
Secondary Storage Disruption

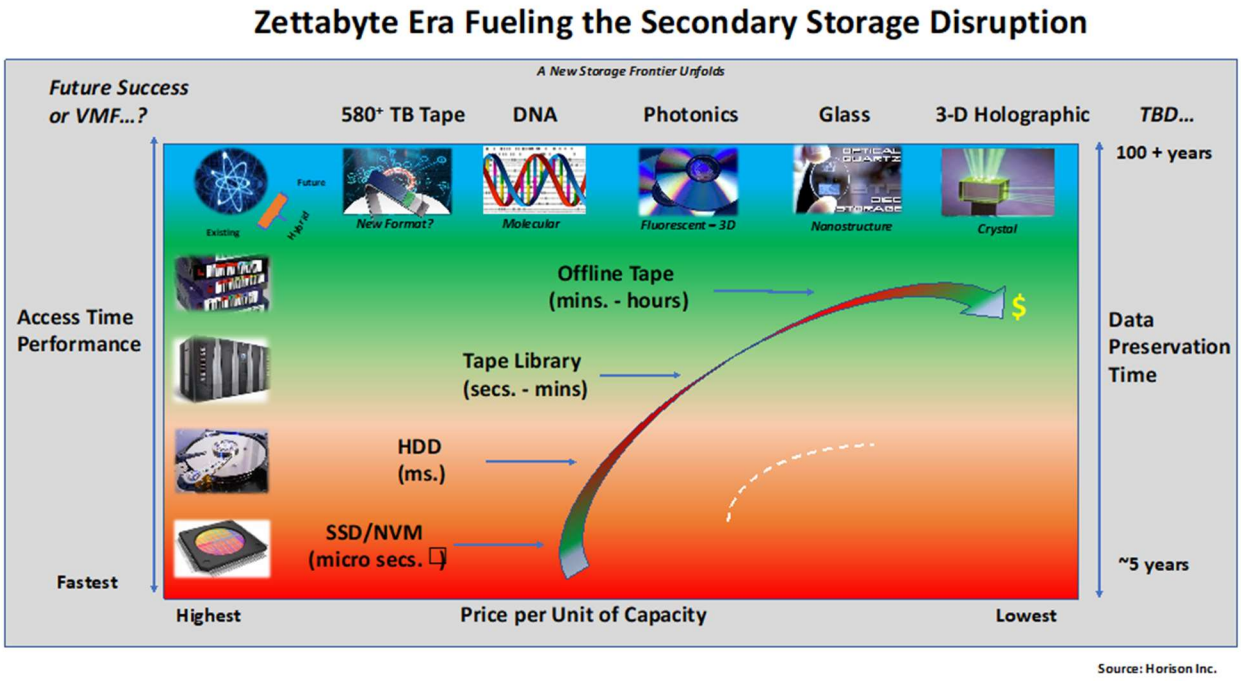

Source: Horison Inc.

Supplier Risk Factors to the Secondary Storage Model

- The zettabyte scale WW secondary storage market is the exclusive domain of few suppliers
- One tape drive manufacturer controlling tape ecosystem (HW and SW) specifications
- Two cartridge tape media suppliers (LTO and enterprise tape formats)
- Three HDD suppliers
- Four primary large-scale robotic tape library suppliers
- Vendors and manufacturers are often dis-incentivized to invest in a novel technology
- IP litigation risks, as evidenced by recent lawsuits, can impact and delay roadmap progress

What Are the Possible Scenarios?

Neither HDD, SSD nor tape technology is currently positioned by itself to effectively meet the enormous HSDC capacity and performance demands, and optical disc is still not a data center class technology. Will the existing tape and HDD roadmaps continue to evolve at current rates, or will they accelerate their development pace? Can a new tape format or cartridge geometry be developed? Could a hybrid combination of existing HDD, tape and SSD technologies be architected? Will any of the brand-new potential technology developments finally break through the labs? Could an independent alliance of HSDCs be formed to define future requirements and pool resources more effectively to accelerate development of the optimal HSDC solution? There are several possible scenarios however secondary storage challenges will mount unabated as we patiently wait for one of them to materialize.

Mr. Moore's Conclusion

Though many enterprise data centers may be able to meet storage demand with current product availability schedules for a while longer, HSDCs are at the VMF epicenter and will soon require a quantum leap or total secondary storage disruption to efficiently meet their long-term demands. The number of HSDCs worldwide projected to reach 1,200 by 2026 adding to the challenge. Without a new trajectory, the secondary storage market is presently left to move forward with incremental investments in existing technologies and will continue falling behind the demand curve. Vendors and manufacturers are often dis-incentivized to invest in a novel technology since the risk reward may not be high enough, while HSDCs leverage their buying and bargaining powers to further drive down prices for existing solutions. The ability of the current secondary storage technology model to meet these demands is in jeopardy as the potential for a VMF increases. The race to zero and few storage manufacturers add more fuel to the potential for a VMF. Several new technology developments are in progress, but their general availability remains unclear. The storage industry has spent the past decade talking about the mounting challenges of extreme data growth – it's time to spend the next decade doing something about it. An impending VMF can be averted.

TRENDFOCUS Analysis

Tape Media Costs – Hyperscale Pricing Trends

Like what has been seen in both HDD and flash-based storage technologies, the mass growth of cloud storage deployment has created what amounts to a completely separate market from traditional enterprise OEM systems sales for on-premises storage usage. With hyperscale customers controlling their own hardware and software stacks and having moved to ODM-direct commodity hardware assembly, hyperscale companies can command market-leading low pricing from their suppliers. The motivation for device suppliers is that the high volumes of a less complex mix of products can support a sustainable long-term business for some companies as hyperscale has become the primary growth driver for enterprise-class storage capacity across all technologies. As Mr. Moore stated in his piece, the race to $0.00 per petabyte challenges all industry players' margins and R&D funding to develop the technology and supply chains to support storage density increases that can be offered at progressively lower dollars per/GB-TB-EB levels.

Tape, up until relatively recently, has maintained a steady demand from on-premises customers who, as purchasers of OEM systems, increased capacity and reduced costs at fairly elongated time intervals. The backwards compatibility of new generations of tape drives to read or write one or two prior generations of media allowed on-premises users to upgrade drives and migrate media to newer, higher capacity versions every couple of generations in many cases. The shift to Ba-Fe media in LTO-7, however, meant that LTO-8 and LTO-9 are only capable of reading and writing cartridges from one generation back (in the case of LTO-8 drives, though, both LTO-7 and the intermediate LTO-7M media can be read and written). Future LTO generations promise to double native capacity through advancements in media, head and drive technologies and will eventually usher in the transition from Ba-Fe media to Sr-Fe alloys. Native capacities from LTO-10 onward may provide the densities required to support hyperscale data growth but may actually exceed storage density for many on-premises customers. Tape companies will likely have to offer different cartridge capacities for non-cloud customers, utilizing shorter tape lengths, resulting in lower capacities suited to the on-premises market.

If tape generations move well beyond the capacity needs of individual on-premises customers – for instance the LTO Ultrium roadmap recently announced that LTO-14 would support up to 576 TB of native capacity – on-premises customers may end up sticking with older generations of LTO drives and media as their capacity requirements are far outstripped by the technology. This type of market split has already been witnessed in the nearline HDD market, where traditional IT systems for on-premises usage skew toward the so-called mid-capacity nearline models, roughly 12 TB or less as of today, while cloud customers progress to 20 TB and beyond. The big difference between HDDs and tape is that HDD storage devices are self-contained, meaning that the matching up of drives and compatible media in the tape market is not an issue for HDDs. However, for the HDD industry, tough decisions must be made regarding whether or not to refresh lower capacity models with newer technologies that can potentially lower costs and preserve margins – for instance, a 2 TB/disk technology can conceivably support lower-cost 12 TB HDDs with 6 disks and without the use of helium sealing used in leading-edge drives. HDD customers, namely system vendors, would be faced with re-qualifying newer generations of HDD models all over again at some significant cost – would the lower pricing of those drives support the additional qualification efforts? Would HDD vendors continue to support the manufacturing of aging generations of lower capacity HDD models as a result? These questions are facing the HDD producers and consumers currently and could very well be facing the tape industry and customers as well in the future.

The architectures and technologies that will help cost-effectively expand the secondary storage market are still waiting to be decided, but clearly, a modest but growing number of hyperscale companies have embraced tape for some of the coldest levels of archival storage for now. As published in prior reports, hyperscale consumption

of tape capacity represented less than one-quarter of total tape capacity shipped in 2021 but is lining up to expand at much faster rates in the coming five years with new customers entering the picture combining with the plan to utilize tape-based services to capture a burgeoning amount of archival data. As a result, compared to past LTO tape generation introductions that took roughly three years for the new generation to cross over the prior one in per-GB costs, the faster uptake by cloud customers will fuel a faster per-GB cross over, dropping to roughly two years in the case of the current LTO-9 transition and as little as one year for LTO-10 which, as of now, should enter the market in 2024. The higher densities of next-generation cartridges will increase media pricing on a per-unit basis but continue to support progressively lower $/GB which, at the system level, reduces overall total cost of ownership (TCO) as more data can be stored within the confines of an existing library. Or, in some cases, the size of a library to store a fixed amount of data can also be reduced, saving on space and power consumption.

Chart 5
Hyperscale Media $/GB by LTO Generation

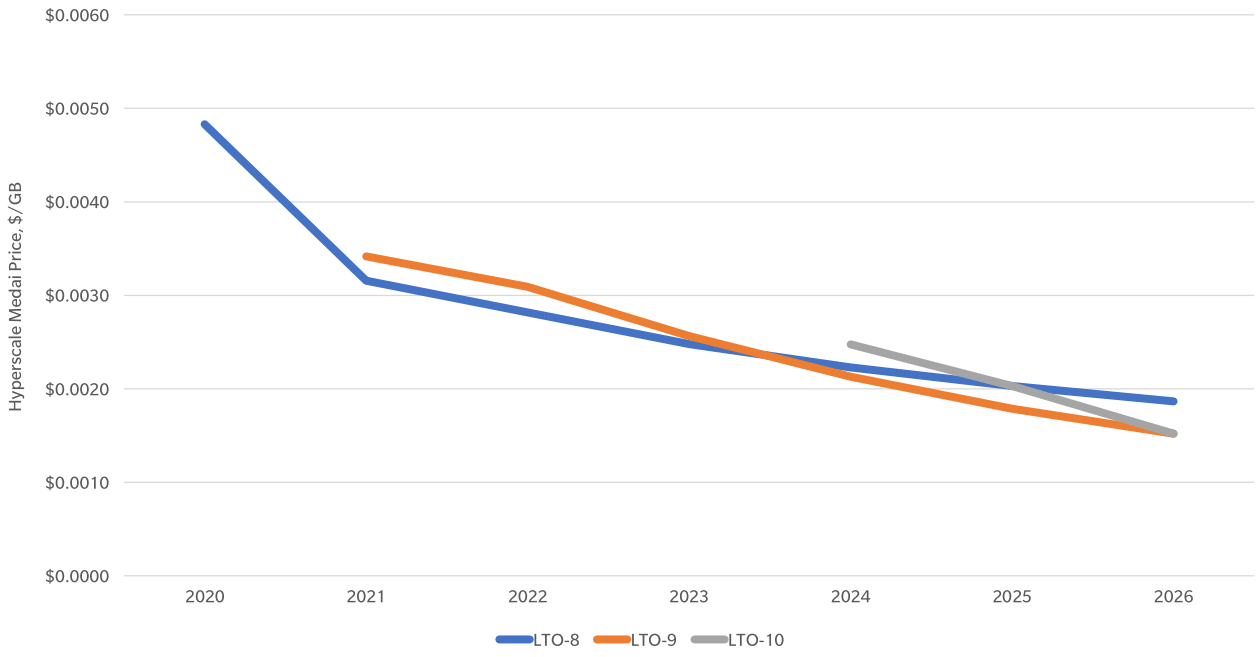

Hyperscale HDD Archive Storage Forecast

In TRENDFOCUS' **Cloud, Hyperscale and Enterprise Storage Service CQ2 '22 Quarterly Update and Revised Long-Term Forecast**, the total nearline unit and exabyte forecast through 2026 was updated to reflect some near-term softening of demand due to weakening economic conditions. While the depth and duration of economic slowing remains uncertain, all TRENDFOCUS forecasts have softened storage demand across most end markets through the first half of 2023 and assume the beginning of recovery to take hold in the second half of next year and continue into 2024. From a long-term exabyte growth CAGR, this near-term slowing and recovery catch-up has little effect on the five-year growth expectation with nearline capacity growth expected to track 29% through 2026 and continue to exceed 3.3 ZB of storage in that year.

Hyperscale capacity growth overall outpaces the demand in OEM and other markets, with a five-year CAGR of 35%, reaching 2.7 ZB of capacity in 2026, or over 80% of all nearline exabytes. Within this category, as was reported in the CQ4 '22 Tape and Archive Storage Service report, around 16% of this capacity was allocated to archival storage services in 2021 or just over 90 exabytes. By 2026, around 20% of hyperscale nearline capacity will be allocated to archive storage services reaching 550 exabytes and tracking a 44% CAGR from 2021 through 2026. This forecast assumes growth in a business-as-usual approach to current archival storage services.

Table 2
Nearline HDD Capacity by End Market, Use Case

	Actual			Revised Long-Term Forecast					
	2019	2020	2021	2022	2023	2024	2025	2026	2021 to 2026 CAGR
Hyperscale Nearline Archive EB	40.65	57.28	90.31	119.55	155.74	248.20	341.75	549.81	44%
Hyperscale Nearline Non-Archive EB	262.38	361.79	485.27	583.67	760.40	1130.68	1556.85	2199.25	35%
OEM/Other Nearline EB	170.46	215.89	352.77	315.49	335.25	467.23	528.90	560.44	10%
Total Nearline EB	473.49	634.96	928.35	1018.71	1251.39	1846.10	2427.50	3309.50	29%

Chart 6
Nearline HDD Capacity by End Market, Use Case

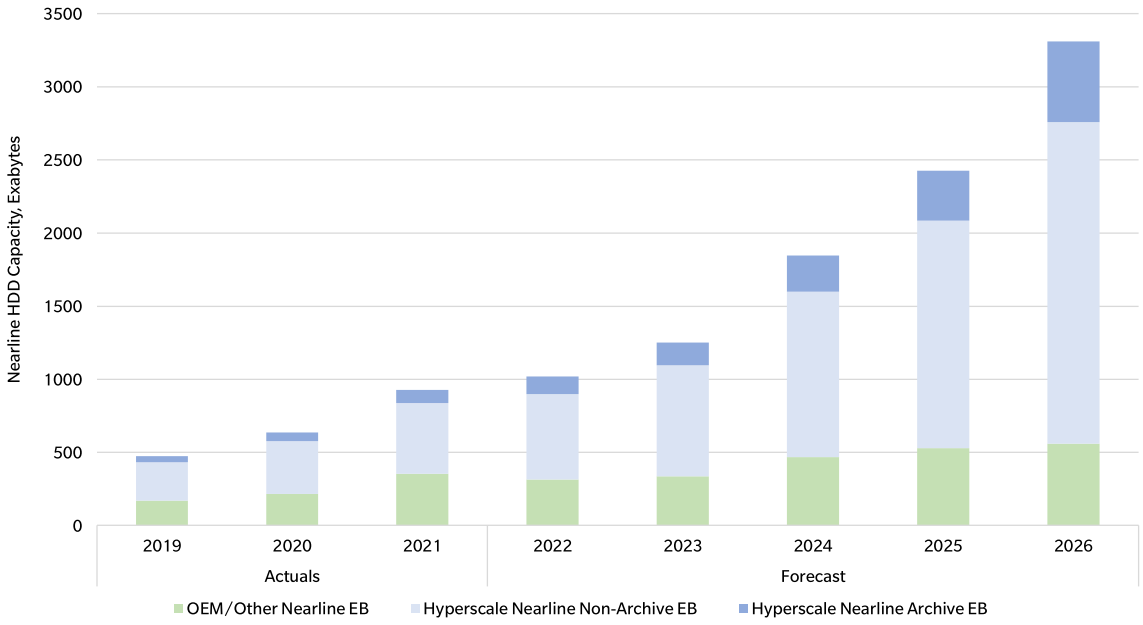

As suggested by Mr. Moore's report, upwards of 80% of stored data can be considered archival in nature, assuming that data can be separated into active archive, archive and deep archive tiers. Arguably most cloud data that would be considered in the active archive and archive tiers today resides in traditional active HDD-based storage services or, as Mr. Moore points out, the "wrong place", whether in the cloud or on-premises. New secondary storage solutions that more tightly integrate multiple storage technologies controlled by intelligent hardware and software will should open opportunities for hyperscale companies and device makers to develop more cost-effective and sustainable storage solutions that preserve valuable data in the appropriate tiers. Over the next decade, it is highly unlikely that any brand-new storage technology other than NAND, HDD and tape will emerge, so evolution and integration of existing technologies will be required. The key for storage providers is to enlarge the percentage of data stored through innovative approaches without cannibalizing current revenue and profit streams required to fund the development secondary storage solutions.