Furthur
Market Research

PRESERVATION OR DELETION: ARCHIVING AND ACCESSING THE DATAVERSE

Inevitably and inescapably, richly varied computing technologies will come and go, ***but the data we create will remain***, and will grow to unimaginable immensity.

This WhitePaper was jointly sponsored by Fujifilm, IBM, and Twist Bioscience.

March 2023

John Monroe | **Furthur Market Research**

TABLE OF CONTENTS

DATA DEFINITIONS:

A byte is a unit of digital information that usually consists of eight bits and is the smallest addressable unit of memory in most computer architectures. A kilobyte (KB) is a thousand bytes of data. A megabyte (MB) is a thousand kilobytes. A gigabyte (GB) is a thousand megabytes. A terabyte (TB) is a thousand gigabytes. A petabyte (PB) is a thousand terabytes. An exabyte (EB) is a thousand petabytes. A zettabyte (ZB) is a thousand exabytes (or a million petabytes). Beyond a zettabyte, the numbers implied by the names become progressively unimaginable (and unpronounceable). A yottabyte (YB)—named after Yoda of Star Wars fame—is a thousand zettabytes. A xenottabyte (XB) is a thousand yottabytes. A shilentnobyte (SB) is a thousand xenottabytes. A domegemegrottebyte (DB) is a thousand shilentnobytes, or 1,000,000,000,000,000,000,000,000,000,000,000 bytes.

NOTE:

We define "enterprise petabytes" as the combined capacities of all enterprise-class hard-disk drives (HDDs), solid-state drives (SSDs), tape, optical, and other potential future building blocks shipped for integration in global enterprise storage systems; the estimated "active installed base" of these enterprise petabytes assumes a 5-year infrastructure replacement/refresh cycle.

HYPERSCALE DATACENTERS:

The largest hyperscale data centers consumed ~65% of all enterprise petabytes delievered in 2022. We believe these companies are currently the top 15 largest global hyperscale cloud service providers, in terms of sheer internal consumption of enterprise petabytes (in alphabetical order): Alibaba, Amazon, Apple, Baidu, eBay/PayPal, Facebook, Google, HUAWEI CLOUD, IBM Cloud, JD.com, Microsoft Azure/ LinkedIn, Oracle Cloud, Salesforce, Tencent, and Yahoo.

PROLOGUE

This WhitePaper is intended to be an extension and refinement of themes explored in our initial WhitePaper —
The Escalating Challenge of Preserving Enterprise Data. *For more details*.

In addition to revised 2023-2030 forecasts based on estimated actual 2022 deliveries of enterprise petabytes (which declined unexpectedly and precipitously during 2H22 for SSDs and HDDs and are still being finalized), we include the survey results of seven interviews with enterprise data center managers and consultants aimed at determining current and future storage needs in the scientific, government public record, corporate, media/entertainment, and video surveillance vertical markets.

Our thesis, antithesis, and inconclusive conclusions remain the same, but we have refined our delineations of the enterprise data layers in our updated forecasts according to access frequency—nanoseconds to years—and estimated the changing size of hot, warm, cool, cold, and frozen data layers.

The Expanding Dataverse

In many guises, and bidirectionally, the IT enterprise has become part of the fabric of our daily lives, and the fabric of our daily lives has become part of the IT enterprise. The constantly growing multidirectional data paths extend from consumers to data centers and back again in exponential ways. Living rooms, PCs, automobiles, mobile phones, and the expanding realms of on-premises and off-premises enterprise storage and computing will increasingly reflect the same global network.

The true value and opportunities generated by this network will be huge but incalculable. Metcalfe's Law states that the "effect" (later revised to "value") of a telecommunications network is proportional to the square of the number of connected users (nodes) of the system. The effective value of multiple billion-node networks interacting chaotically and unpredictably is immeasurable with any degree of precision.

But what is certain is that the people and systems and sensors connected in this network have generated and will continue to generate immense quantities of data. It has now become a cliché to say that "data is the new oil"—unlike oil never to be burned but like oil always to be mined for its potential value. It should be noted that clichés earn their status as clichés because they are so obviously true.

The enterprise "data pools" of the early 2000s became "data lakes" by 2010 and grew to become "data oceans" by 2015 and have already begun morph into a vast multiform "dataverse." The active installed base of enterprise petabytes exceeded 1 zettabyte (a million petabytes) in 2016 and grew to 5.2 zettabytes (5.2 million petabytes) in 2022.

We are squarely in the midst of the "zettabyte era." And we are loathe to delete any data.

Thesis

Despite wider deployments of storage efficiency technologies, demand for enterprise storage capacity delivered on various forms of media will continue to expand at explosive rates—at least 30% to 35% per year—through 2030 and perhaps much longer. This means there likely could be a real need for application-centric enterprise systems to store, secure, manage, back up, analyze and derive advanced, intelligent value from ~13 million freshly delivered petabytes of mostly unstructured enterprise data capacity in 2030, bringing the total active installed base of enterprise data to ~39 million petabytes, up from approximately 91,000 petabytes in 2010.

Antithesis

Although the amount of freshly shipped enterprise data capacity expanded at an actual compound annual growth rate (CAGR) of ~30% from 2010 through 2022, budgetary constraints combined with new AI/ML-enabled data reduction technologies—and with evolving and more strictly administered archive "rules," and with limited availability of enterprise-grade media—will diminish this rate of increase from 2022 through 2030 and perhaps much longer. There will be an increasingly large gap between the data that is generated by billions of sensors and devices and systems and the data that is actually stored to be mined for current and future value on enterprise-grade media. Diverse inhibiting technological and financial factors combined with greater asset utilization and limited IT budgets will curtail 2022 through 2030 growth rates to ~25%. In a 25%-per-annum-rate-of-increase scenario, there will be only ~8 million freshly delivered petabytes of mostly unstructured enterprise data capacity in 2030, bringing the total active installed base of enterprise data to a potentially far more manageable yet still immense ~26 million petabytes, up from approximately 91,000 petabytes in 2010.

FORECAST SCENARIOS

Table 1 details estimates of enterprise SSD, HDD and tape shipments in petabytes, and provides percentages of total shipments for SSDs, HDDs, and tape, and also depicts active installed base estimates of total enterprise petabytes, 2010-2030.

Table 1: Estimated Enterprise SSD, HDD and Tape Shipments and Enterprise Active Installed Base Estimates, 2010-2030

	2010	2015	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	"CAGR 2023-2030"
Enterprise SSD PB	187	26,154	130,766	178,972	186,588	230,664	334,195	498,792	690,849	925,047	1,098,030	1,604,222	2,050,196	36.6
Annual Growth %	–	223.9	64.1	36.9	4.3	23.6	44.9	49.3	38.5	33.9	18.7	46.1	27.8	
Enterprise HDD PB	45,216	157,093	679,887	959,011	941,749	1,251,584	1,610,789	2,177,786	2,735,300	3,605,125	4,589,324	6,021,194	8,062,378	30.5
Annual Growth %	–	34.7	39.9	41.1	−1.8	32.9	28.7	35.2	25.6	31.8	27.3	31.2	33.9	
Enterprise Tape PB	30,208	98,432	136,119	189,938	206,842	259,794	303,440	366,252	464,773	574,460	763,457	966,537	1,258,431	25.3
Annual Growth %	–	15.0	1.4	39.5	8.9	25.6	16.8	20.7	26.9	23.6	32.9	26.6	30.2	
Total Compressed Shipments PB	75,611	281,679	946,772	1,327,921	1,335,179	1,742,043	2,248,424	3,042,830	3,890,922	5,104,631	6,450,812	8,591,952	11,371,005	30.7
Annual Growth %	–	33.9	33.6	40.3	0.5	30.5	29.1	35.3	27.9	31.2	26.4	33.2	32.3	
SSD PB % of Total Annual Shipments	0.2	9.3	13.8	13.5	14.0	13.2	14.9	16.4	17.8	18.1	17.0	18.7	18.0	
HDD PB % of Total Annual Shipments	59.8	55.8	71.8	72.2	70.5	71.8	71.6	71.6	70.3	70.6	71.1	70.1	70.9	
Tape PB % of Total Annual Shipments	40.0	34.9	14.4	14.3	15.5	14.9	13.5	12.0	11.9	11.3	11.8	11.2	11.1	
Active Installed Base PB	**91,000**	**819,949**	**2,923,201**	**3,950,945**	**5,232,405**	**6,447,587**	**7,985,007**	**10,081,065**	**12,644,067**	**16,413,518**	**21,122,288**	**27,465,816**	**35,793,990**	

HDD capacities are raw/uncompressed, since so few enterprise HDDs utilize any form of data compression. We assume a conservative average annual expansion rate of ~30% 2023-2030.

SSD capacities reflect an approximate 5x compression ratio, but only for approximately 5% of all enterprise SSD PBs shipped, the vast majority of which (currently ~95%) are configured in server/direct-attached storage (DAS) systems, with little or no data compression, not in fabric-attached solid-state arrays (SSAs), wherein sophisticated data compression software is the norm.

Tape capacities reflect a global average of 2.5x data compression.

Enterprise optical shipments remain minimal at ~1,500 petabytes per year—less than half of 1% of the 2022 total—and have not been included in our estimates of historical shipments, forecasts, or the active installed base.

That said, there surely could be room for enterprise optical to play a strategic role in future enterprise storage markets.

We estimate the active installed base of total enterprise capacities was 91 exabytes in 2010 and will likely grow to ~35 zettabytes in 2030.

For the active installed base, we assume a 5-year infrastructure refresh/replacement cycle, retiring all 2010 shipments in 2015 while adding 2015 shipments to the 2014 installed base, and we repeat this cycle through 2030.

Source: Furthur Market Research (March 2023)

Table 2 depicts alternate 25%, 35%, and 45% growth scenarios as well as various estimates of the active installed base of total enterprise petabytes, 2023-2030.

Table 2: Alternate Shipment and Installed Base Estimates, 2023-2030

	2023	2024	2025	2026	2027	2028	2029	2030
Alternate 2023-2030 Shipment Scenarios								
Total Shipped Enterprise PB Expanding at 25%/Year 2023-2030	1,668,974	2,086,218	2,607,772	3,259,716	4,074,644	5,093,306	6,366,632	7,958,290
Total Shipped Enterprise PB Expanding at 35%/Year 2023-2030	1,802,492	2,433,365	3,285,042	4,434,807	5,986,990	8,082,436	10,911,288	14,730,239
Total Shipped Enterprise PB Expanding at 45%/Year 2023-2030	1,936,010	2,807,215	4,070,462	5,902,169	8,558,145	12,409,311	17,993,501	26,090,576
Alternate 2023-2030 Active Installed Base Scenarios								
Active Installed Base PB at 25% Annual Shipment Expansion	6,374,518	7,749,733	9,410,733	11,342,528	14,081,993	17,433,256	21,551,464	26,466,924
Active Installed Base PB at 35% Annual Shipment Expansion	6,508,036	8,230,398	10,568,668	13,675,554	18,327,364	24,667,757	33,330,622	45,018,031
Active Installed Base PB at 45% Annual Shipment Expansion	6,641,554	8,737,766	11,861,455	16,435,704	23,658,670	34,325,938	50,071,015	73,118,761

Source: Furthur Market Research (March 2023).

Figure 1: Alternate 2023-2030 Active Installed Base Scenarios

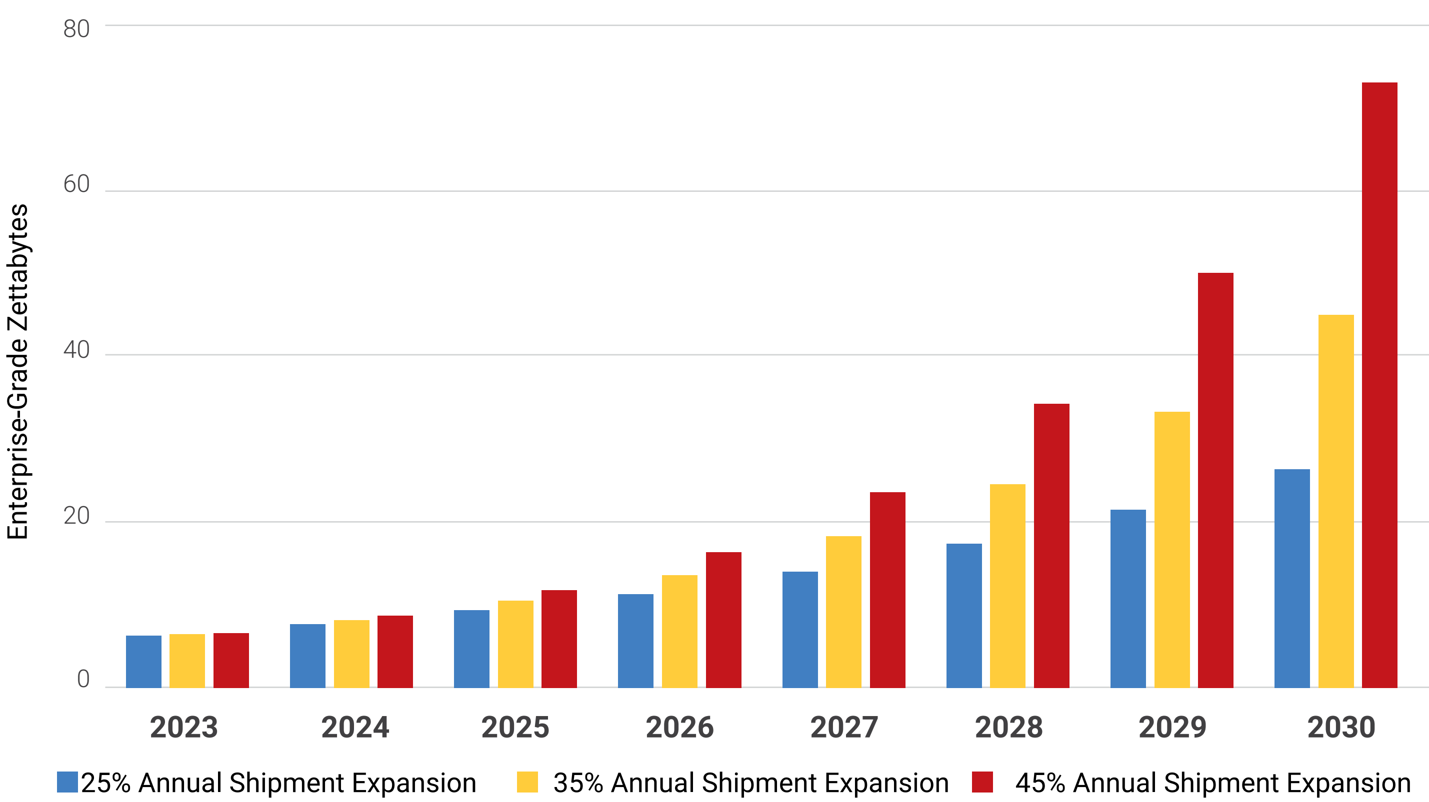

Source: Furthur Market Research (March 2023).

Table 3 details new hot, warm, cool, cold, and frozen delineations of enterprise data "temperature" based on frequency of access, 2010-2030.

Table 3: Estimated Data Temperature and Access Frequencies of Evolving Shipments, 2010-2030

2020-2030 Percentages	2010	2015	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030
Hot Data 8% (nanoseconds to milliseconds) of Shipped PB	6,049	22,534	75,742	106,234	106,814	139,363	179,874	243,426	311,274	408,371	516,065	687,356	909,680
Warm Data 17% (milliseconds to seconds) of Shipped PB	12,854	47,885	160,951	225,747	226,981	296,147	382,232	517,281	661,457	867,787	1,096,638	1,460,632	1,933,071
Cool Data 19%-14% (minutes to 24 hours) of Shipped PB	18,147	61,124	179,887	250,977	238,997	306,599	379,984	495,981	618,657	791,218	941,818	1,228,649	1,569,199
Cold Data 35%-26% (days to weeks) of Shipped PB	28,732	104,503	331,370	463,444	463,307	601,005	773,458	937,192	1,159,495	1,429,297	1,715,916	2,233,908	2,956,461
"Frozen" Data 21%-35% (weeks to years) of Shipped PB	7,561	45,632	198,822	281,519	299,080	398,928	532,876	848,950	1,140,040	1,607,959	2,180,374	2,981,407	4,002,594
Total Shipments as Forecast in Table 1 (~30.7% Annual Growth)	73,343	281,679	946,772	1,327,921	1,335,179	1,742,043	2,248,424	3,042,830	3,890,922	5,104,631	6,450,812	8,591,952	11,371,005
PB Shipped to the Cold/"Frozen" Data Layers	36,293	150,135	530,193	744,964	762,387	999,932	1,306,334	1,786,141	2,299,535	3,037,256	3,896,290	5,215,315	6,959,055
Cold/"Frozen" Data Layer %	49.5	53.3	56.0	56.1	57.1	57.4	58.1	58.7	59.1	59.5	60.4	60.7	61.2
Vendor Revenue Opportunity In Billions of USD						$5.1	$5.1	$6.0	$6.4	$7.0	$7.4	$8.2	$9.0
Enterprise Tape PB Shipments	30,208	98,432	136,119	189,938	206,842	259,794	303,440	366,252	464,773	574,460	763,457	966,537	1,258,431
% of Total Shipments	**41.2**	**34.9**	**14.4**	**14.3**	**15.5**	**14.9**	**13.5**	**12.0**	**11.9**	**11.3**	**11.8**	**11.2**	**11.1**

Source: Furthur Market Research (March 2023)

Figure 2: **The Evolving Storage Pyramid**

Differing degrees of storage temperature, differing technologies in the layers... With an ever-increasing base of cold/frozen data...

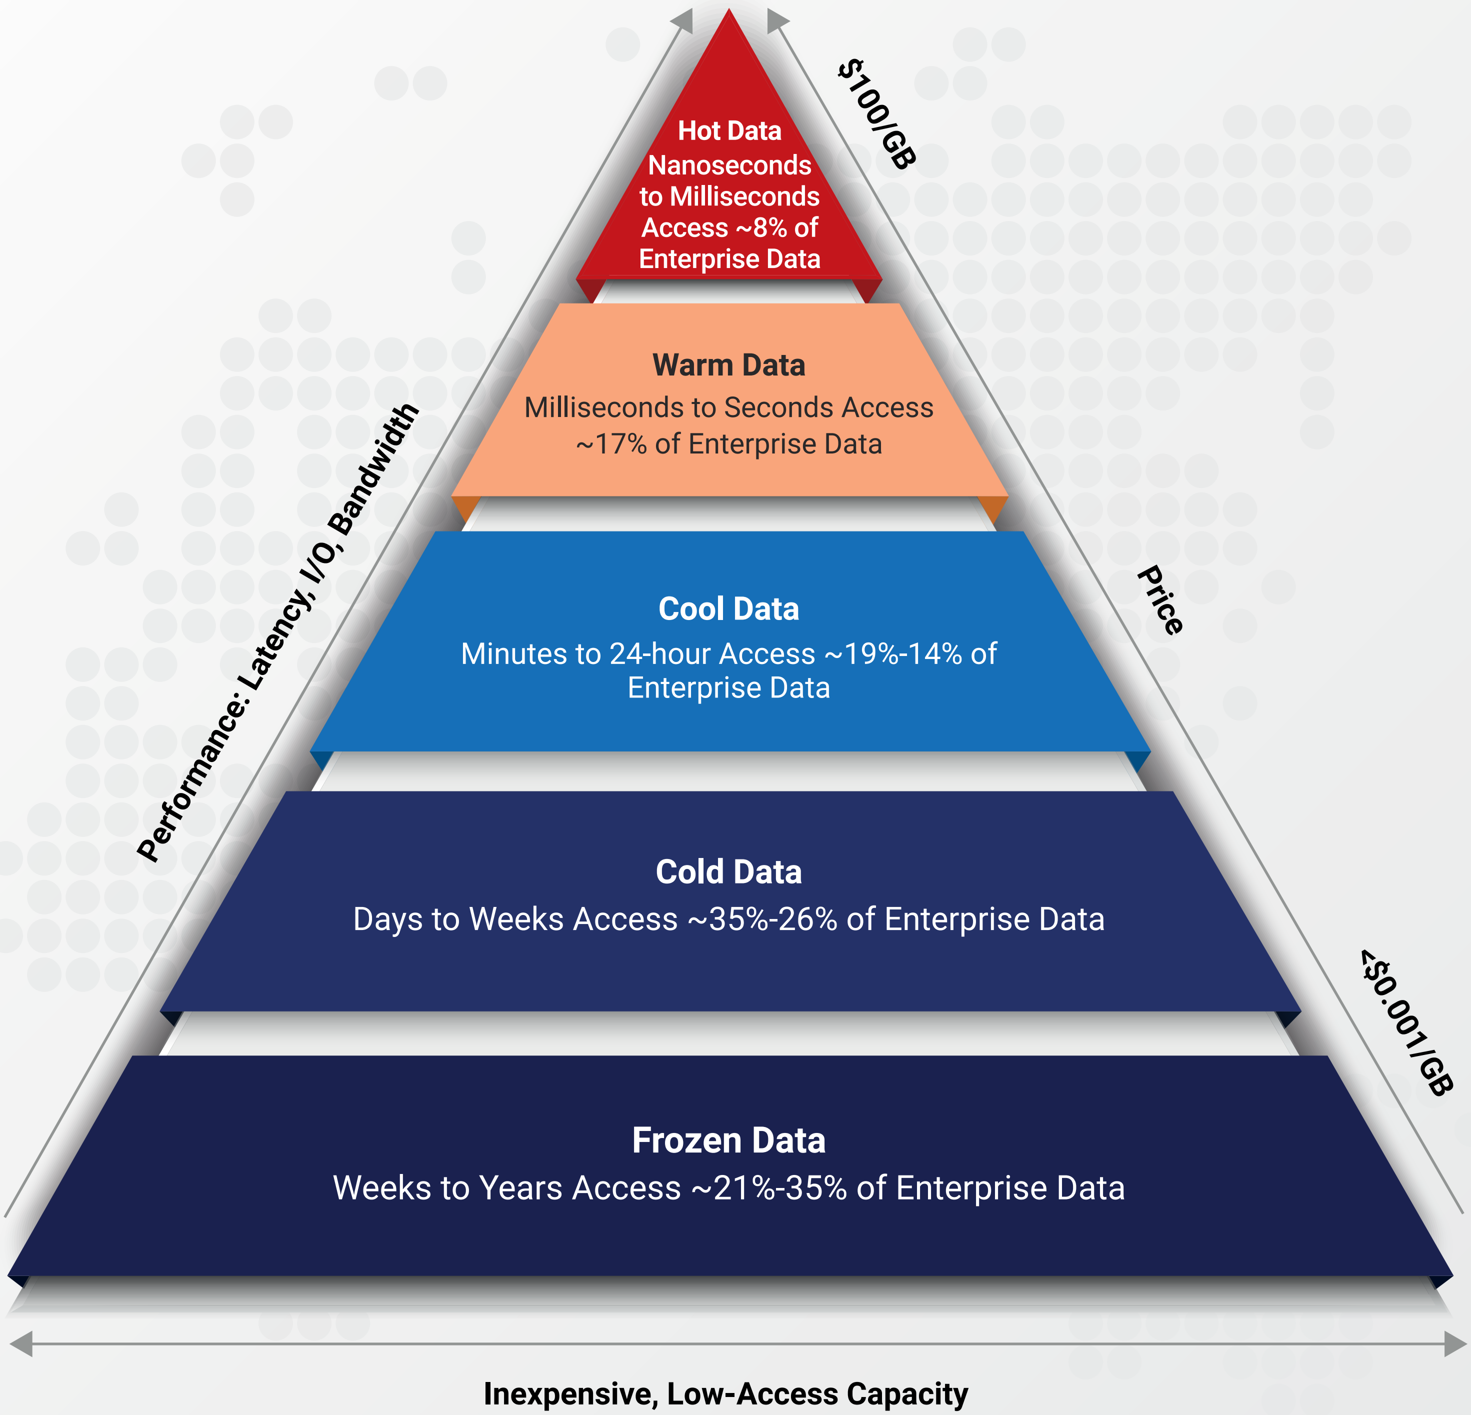

A majority of the cold/frozen layers may be JIC (Just in Case) or WORN (Write Once Read Never) data, which may never be accessed at all–nor, in most cases, will it ever be deleted.

Source: Furthur Market Research (March 2023).

Table 4 portrays revenue opportunities that might be achieved by delivering enterprise storage products at aggressive best ex-factory ASPs per gigabyte designed to manage only the cold/frozen data layers, 2023-2030.

Table 4: Estimated Revenue Opportunities Generated by Managing the Cold/Frozen Data Layers, 2023-2030

	2023	2024	2025	2026	2027	2028	2029	2030
Total Shipments as Forecast in Table 1 (30.7% Annual Growth)								
PB Shipped to the Cold/Frozen Data Layers	999,932	1,306,334	1,786,141	2,299,535	3,037,256	3,896,290	5,215,315	6,959,055
Cold/Frozen Data Layer %	57.4	58.1	58.7	59.1	59.5	60.4	60.7	61.2
Vendor Revenue Opportunity In Billions of USD	$5.1	$5.1	$6.0	$6.4	$7.0	$7.4	$8.2	$9.0
Alternate 2023-2030 Shipment Scenarios								
Total Shipped Enterprise PB Expanding at 25%/Year 2023-2030	1,668,974	2,086,218	2,607,772	3,259,716	4,074,644	5,093,306	6,366,632	7,958,290
Total Shipped Enterprise PB Expanding at 35%/Year 2023-2030	1,802,492	2,433,365	3,285,042	4,434,807	5,986,990	8,082,436	10,911,288	14,730,239
Total Shipped Enterprise PB Expanding at 45%/Year 2023-2030	1,936,010	2,807,215	4,070,462	5,902,169	8,558,145	12,409,311	17,993,501	26,090,576
Evolving Opportunities to Manage Cold/"Frozen" Data, Alternate Scenarios								
57.4%-61.2% ofTotal Shipped PB Expanding at 25%/Year 2023-2030	957,991	1,212,093	1,530,762	1,926,492	2,424,413	3,050,890	3,826,346	4,790,891
Vendor Revenue Opportunity In Billions of USD	$4.9	$4.7	$5.1	$5.4	$5.6	$5.8	$6.0	$6.2
57.4%-61.2% ofTotal Shipped PB Expanding at 35%/Year 2023-2030	1,034,631	1,413,785	1,928,320	2,620,971	3,562,259	4,841,379	6,557,684	8,867,604
Vendor Revenue Opportunity In Billions of USD	$5.3	$5.5	$6.4	$7.3	$8.2	$9.2	$10.3	$11.5
57.4%-61.2% ofTotal Shipped PB Expanding at 45%/Year 2023-2030	1,111,270	1,630,992	2,389,361	3,488,182	5,092,097	7,433,177	10,814,094	15,706,527
Vendor Revenue Opportunity In Billions of USD	$5.7	$6.4	$8.0	$9.7	$11.7	$14.0	$17.0	$20.4
Aggressive Best Ex-Factory ASPs/GB for Tape or Other Technology	**$0.0051**	**$0.0039**	**$0.0033**	**$0.0028**	**$0.0023**	**$0.0019**	**$0.0016**	**$0.0013**

Source: Furthur Market Research (March 2023)

We consider the overall 30.7% 2023-2030 annual growth rates depicted in Table 1 to be conservative. We saw historical ~30% average rates of actual annual growth from 2010 to 2022, and we foresee good chances for the average annual rates of 2023-2030 data demand to expand beyond 30%—as they did in 2014 (up 34.6%), 2015 (up 34.0%), 2019 (up 34.5%), 2020 (up 35.3%), and 2021 (up 40.0%).

Growth rates will surely fluctuate, but we do not think the unprecedented and probably anomalous <10% growth rates we saw in 2022—with hyperscale and other large enterprise customers delaying orders to an inordinate degree and digesting inventory during 3Q22 and 4Q22—will be repeated with any cyclic frequency 2023-2030.

We believe a 25% per annum growth rate in enterprise petabytes delivered is a likely “worst case” scenario while a 35% per annum growth rate is a “likely case” scenario and a 45% per annum growth rate is a perhaps (un)likely “best case” scenario, but it should at least be considered.

In the ~30.7% growth scenario depicted in Table 1, the active installed base will grow to 35.8 zettabytes in 2030. In alternate scenarios, the active installed base of enterprise petabytes will grow from ~5.2 million petabytes in 2022, to ~26.5, ~45.0, or ~73.1 million petabytes in 2030, respectively, in 25%, 35%, or 45% annual growth rates.

Assuming the cold and frozen data layers will likely expand to 61.2% of total shipments in 2030, up from 57.4% in 2023, these layers will comprise ~5, ~7, ~9 or 16 zettabytes in 2030 in 25%, 30.7%, 35%, or 40% growth scenarios.

It should be noted that a majority of the cold/frozen layers may be JIC (Just In Case) or WORN (Write Once Read Never) data, which may never be accessed at all—nor in most cases will it ever be deleted.

Several enterprise IT managers with whom we spoke—but who were not part of this survey group—stated that an exacerbated problem with data deletion is establishing generally agreed-upon ground rules. When they asked for buy-in from their internal clients, they could not obtain any solid commitment for, say, 5-year, 7-year, or 10-year deletion objectives for aging data. There was always the lingering fear that after 5 years or 7 years or 10 years and 1 day, they would absolutely need that old data for some unspecified, but critical, future purpose.

THE GROWING ENORMITY OF FROSTBITTEN DATA

In Table 3 and Figure 2, our assumption is that the hot and warm data layers will remain fairly constant at ~25% of the total (8% hot, 17% warm). The cool data layer shrinks somewhat from ~19% in 2020 to ~14% in 2030 and the cold data layer shrinks as well, from 35% in 2020 to 26% in 2030, while the frozen data layer grows from ~21% in 2020 to ~35% in 2030.

With the advent of more strictly enforced AI/ML corporate archive and access rules, it is also possible that, in many data centers, 25% of the data will grow to be hot in 2030, while the warm and cool data layers decline to insignificance, and the cold/frozen data layers will grow to 75% of the total.

The active installed base of all enterprise petabytes exceeded one zettabyte in 2016 (the beginning of the “zettabyte era”). We estimate that new shipments of enterprise storage capacity only to the cold/frozen data layers will exceed one zettabyte in 2023. By the end of 2023, it is likely that the cold/frozen data layers—with access frequencies spanning days to years—will comprise 3.7 zettabytes (~57%) of the total 6.5 zettabytes in the active installed base.

Conservatively based on our aggressive best estimated ASPs/GB (which may not be feasible or profitable, and will be available only to the largest enterprise customers), annual vendor revenue opportunities derived from managing only the cold/frozen data layers could expand to ~$6.2, ~$11.5, or ~$20.4 billion, respectively, in 25%, 35%, or 45% annual expansion scenarios. Industry average rather than best ASPs/GB should increase these vendor revenue opportunities by at least 10%, and these vendor revenue estimates should increase by at least 25% if we consider end-user markups for integrated storage systems rather than merely for the essential storage building blocks.

SUSTAINABILITY

In 2022, we estimate that the cold and frozen data layers accounted for 762.4K of the total enterprise petabytes delivered (and, as mentioned above, will account for more than one zettabyte of the enterprise capacity delivered in 2023), while tape accounted for only 206.8K of these petabytes in 2022 (see Tables 1 and 3).

It is obvious that HDDs and perhaps a significant number of SSDs are handling far too much of the cold/frozen workloads at far too great a cost/GB while consuming an inordinate share of available energy.

Table 5 estimates the total power consumption of enterprise HDDs, SSDs, and tape from 2020-2025.

Table 5: Enterprise HDD, SSD, and Tape Power Consumption Estimates, 2020-2025

Total Shipments, Interfaces	Active/Idle Watts	% of Shipments	% Active/Idle Usage	Total Power Watts		Total Power Megawatts
Enterprise-Grade HDDs ~428M Shipped 2020-2025						
SATA Idle	5.5W	85%=364M	65	1,301,300,000		1,301
SATA Active	7.7W		35	980,980,000		981
SAS Idle	5.8W	15%=64M	65	24,128,000		24
SAS Active	9.8W		35	21,952,000		22
					Total for HDDs	**2,328**
Enterprise-Grade SSDs ~399M Shipped 2020-2025						
SATA/SAS Idle	1.5W	20%=90M	40	48,000,000		48
SATA/SAS Active Read/Write	2.1W/3.2W		60	100,800,000		101
NVMe Idle	3.5W	80%=319M	40	446,600,000		447
NVMe Active Read/Write	11W/13.5W		60	2,105,400,000		2,105
					Total for SSDs	**2,701**
HDD+SSD Power Consumption vs Enterprise Tape						
Estimated Total 2020-2025 Megawatt Power Consumption for New Shipments of HDDs + SSDs						**5,029**
Estimated Total 2020-2025 Megawatt Power Consumption for HDDs + SSDs in the Active Installed Base (~3x New Shipments)						**15,087**
Estimated Total 2020-2025 Megawatt Power Consumption for Enterprise Tape in the Active Installed Base						**18**
					Ratio HDD+SSD:Tape	**838**

Source: *Furthur Market Research (March 2023)*

Contrary to popular belief, we estimate that, with the advent of PCIe interfaces, enterprise SSDs will consume ~16% more power than enterprise HDDs during the 2020-2025 time period. However, higher-performance SSDs can reduce the need for additional servers, which was a central concern of our corporate interview (see Table 8), but this will be a complexly variable cost saving to calculate, depending on the data center infrastructure. Regardless of substantial rack and cooling and other infrastructure costs, just powering these essential storage building blocks will require ~5,029 megawatts.

If we try to calculate this for the active installed base, assuming an approximate 3:1 uptick for total units in the active installed base, ~15,087 megawatts of power will be consumed 2020 through 2025 just to fuel the online activities of enterprise HDDs and SSDs—a stunning 838 times more than tape.

> Limited SSD and HDD production capabilities measured against actual storage and sustainability needs may drastically alter the potential market size and demand for new generations of tape, DNA data storage, and optical technologies.

The need for ultra-low-cost, sustainable storage alternatives to manage the cold/frozen data layers is blindingly blatant. At the moment, enterprise tape is the only viable alternative, and the hyperscalers have recently recognized the virtues of tape and have become the largest customers, eclipsing by orders of magnitude tape demand from traditional enterprise accounts.

Despite hugely increasing shipments to global hyperscale accounts, tape as a percentage of total enterprise petabytes delivered has declined drastically in recent years, dropping from 34.9% in 2015 to 15.5% in 2022, and may continue to decline to 11.1% in 2030—even while expanding greatly from 206.8K petabytes to 1.3 zettabytes. That said, we believe there will be a resurgence in tape shipments for a variety of reasons based on expanding demand on multiple fronts, relative data temperature and time-to-data needs based on access frequency, and lower costs of data retention and power consumption, as well as limited HDD and SSD production capabilities. Tape could well grow to at least two zettabytes delivered in 2030.

Global demand may certainly swell to ~11-15 million enterprise petabytes per year 2028-2030. We have forecast HDD and SSD shipments to expand to meet this demand, but we have deepening doubts about maximum production capabilities, partially because the costs and availability of the basic building blocks of enterprise storage will undergo diverse transformations during the next decade. Though storage prices to customers will continue to decline on a per-bit basis, the expense to produce the bare bits of advanced technologies on enterprise-grade media—and resale prices to customers—will decline at slower rates during the 2023-2030 than in the 2010-2022 time period, and maximum available capacities may be limited.

Tempered by a colorful history of profitless price wars caused by needless surplus production—which, even after the 2012 consolidations, wherein Seagate acquired Samsung's HDD division and Western Digital acquired Hitachi GST, sadly recurred with a regularity that rivaled the seasons—the HDD makers have grave fiscal concerns about investing unprofitably in future CAPEX in the face of uncertain demand and growing SSD incursions. We fear the HDD industry will not adequately invest to be able to deliver ~5 million, much less ~8 million, enterprise-grade petabytes per year from 2028-2030. And given the recent precipitous price erosions—the price for raw NAND dropped by more than 70% during 2H22—and the inevitability of future supply/demand imbalances and the attendant price fluctuations, we also have growing doubts that the NAND industry will spend the necessary hundreds of billions of dollars to be able to deliver ~1 million, much less ~2-3 million, enterprise-grade SSD petabytes per year from 2028 to 2030.

Limited SSD and HDD production capabilities measured against actual storage and sustainability needs may drastically alter the potential market size and demand for new generations of tape, DNA data storage, and optical technologies.

The LTO and IBM Jaguar/Enterprise Tape roadmaps represent a faster path to greater per-drive/cartridge storage capacities than HDDs—LTO 9 is already delivering 45TB compressed per cartridge, versus 20-22TB uncompressed for enterprise HDDs, and should certainly deliver >300TB compressed per cartridge prior to 2030, whereas we believe it is unlikely that enterprise HDD per-drive capacity will exceed 50TB prior to 2030—and the tape suppliers can surely ramp more rapidly and cost-effectively to greater petabyte shipments than HDDs or SSDs simply by shipping more cartridges to expand the capacity of new and installed libraries. There will be an increasingly immense hunger for the least-expensive ASPs/GB combined with the greatest available capacities, and successive generations of tape or other inexpensive enterprise storage technologies should play a crucial role here.

Enterprise tape, in its current and future configurations, not only will have substantially lower initial acquisition costs/GB than enterprise HDDs or SSDs—as of 4Q22, the best available OEM prices were $0.0731/GB for SSDs (unprofitable), $0.0096/GB for HDDs (marginally profitable), and $0.0051/GB for tape (profitable), which we forecast will decline to ~$0.0179/GB for SSDs, ~$0.0061/GB for HDDs, and ~$0.0013/GB for tape or other technology in 2030—but will also consume far less power, since the vast majority of tape cartridges will not at any given moment actually be mounted in the tape drive itself, but reside offline, either in a slot of the robotic library, or in some off-premises location, effectively drawing minimal or no power. Based on current forecasts (and far fewer tape petabytes deployed), the active installed base of enterprise tape will likely consume only ~18 megawatts of power from 2020 through 2025, a staggering 838 times less than the active installed base of HDDs and SSDs, and, unlike HDDs and SSDs, will be largely impervious to malware attacks, because of "offline air-gap" protection of sensitive enterprise data.

As global storage needs expand 2025-2030, other low-cost, low-power technologies such as DNA data storage or new breeds of enterprise optical will surely be required to service developing demand.

SURVEY RESULTS

We interviewed via phone and email seven managers of IT storage technologies which gathered and preserved data from the following vertical markets:

In the case of **video surveillance**, our interview was with a consultant who had in-depth knowledge of the needs of diverse organizations dealing with data from **500 to 2,000** cameras operating **24/7/365**, such as hospitals, universities, casinos, airports, sports arenas, and theme parks.

Tables 6, 7, and 8 summarize answers given to our posed questions.

Table 6: Survey Results 1

Vertical Market	Multiple or Single Data Center(s)	%On-Prem/Off-Prem; %Edge/Hybrid	Database Size; Primary Application; Number of Users Serviced	%Hot/Warm/Cool/Cold/Frozen
Scientific 1	One main, but three others "borrowed" globally for backup support	100% on-prem, backup support DCs 100% on-prem, 0% "cloud" storage; 0% Edge/Hybrid	500PB; on-site scientific experiments; thousands of users	100% "cold," but can become "hot" at any time depending on data access requests; HDD caches 200PB useable, and 10s of PBs of SSD between HDD and tape; initial "data gathering" 100% to tape at 40GB/sec
Scientific 2	Three, geographically dispersed	100% on-prem and 100% backed up on-prem, but also 100% backed up to cloud, with a guarantee that all data is in three different physical locations, with 4-6-hour retrieval times; 0% Edge/Hybrid	100PB total capacity, currenlty 22PB used; raw data gathered 24/7 from myriad worldwide sources; thousands of users	100% "cold," but can become "hot" (~8%) at any time depending on data access requests; all data could be considered JIC but none of the data is WORN; assumption is that all data will be accessed eventually
Scientific 3	Two, both on-prem, one new as of 4Q22	100% on-prem, but part of the data center is used to provide "off-prem" backup support for another distant site; 0% Edge/Hybrid	212PB with 4PB (<2%) HDD cache; on-site scientific experiments; thousands of users	100% "cold," but can become "hot" at any time depending on data access requests; initial "data gathering" 100% to tape; no further delineation or tiering will be necessary, the application remains the same irrespective of expanding database size
Government Public Records	Multiple, geographically dispersed	"Enterprise" (active) 90% on-prem, 10% off-prem "Preservation" (archive) 2 copies on-prem, 1 copy off-prem; 0% Edge/Hybrid	"Enterprise" 59PB, Preservaton" 25PB; government and all other published records; potentially millions of users	"Enterprise" 100% Hot "Preservation" Combination of Warm/Cold, mostly Cold/Frozen, but almost never WORN
Media/ Entertainment	4 principal sites: 2 archive off-site "mirrors," 2 for data ingestion, production, and editing	100% on-prem, corporate owned and managed; a small off-prem cache for distribution of certain products/services is insignificant; 0% Edge/ Hybrid	~110PB; real-time video feeds from multiple sources; potentially tens of millions of viewers per event	~14PB "hot," live-feed edited/ re-compiled in real-time; all data becomes "cold" but any portion this active archive may become relevant and "hot" in real-time; the tape data centers are a ~20-minute drive from the "live" production/editing centers, any portion of the archive can be accessed in ~20 minutes
Corporate	10 multi-function data centers, geographically dispersed	100% on-prem, corporate owned and managed; 0% Edge/Hybrid	~50PB; a proprietary system that combines many scientific and corporate applications; thousands of users	100% "cold" but 80% of this archive is actively accessed
Large Video Surveillance Customers (with 500 or more cameras)	Usually only one, on-site	100% on-site; 0% Edge/ Hybrid	Per camera per day: 109.4GB 500 cameras: 4.9PB 90-day retention 1000 cameras: 9.8PB 90-day retention 2000 cameras: 19.7PB 90-day retention Usually accessed by less than a dozen IT analysts	100% "cold" but could become "hot" at any time searching for video proof of break-in/burglary, injury claims, medical malfeasance, employee theft, suspicious baggage /behavior, card counting, student assault, etc

Source: Furthur Market Research (March 2023)

Table 7: Survey Results 2

Vertical Market	Annual Database Expansion Rate: Retention Period of Immutable Data	Optimal Size of Data Objects	Infrastructure Refresh Time Periods	Structured vs Unstructured Data %
Scientific 1	25% somewhat variable, but stable long-term; Indefinite	Combine myriad MB-size files, 100GB optimal	HDD/SSD ~5 years, tape in concert with technology evolution ~3-6 years; library robotics ~10 years	100% structured in that the physical experiments are "structured"
Scientific 2	20%, expanding; Indefinite	Small file streams aggregated to 250K to 25GB; ~750K files/ day; too large a file gums up the system, all files are backedup 100% on-prem and off-prem simultaneously	SSD ~3-5 years, HDD ~5-7 years, tape drives ~5 years, tape robotics ~15 years	100% structured, with sophisticated and extensive metadata to create easy access to specific data sets
Scientific 3	10% now but soon expanding to 25%; Indefinite	2GB to 4GB; anything smaller causes wasted seek times	~3-5 years, in concert with advancing LTO tape technology	~65% "raw/unstructured" but file optimization narrows down the need to store 100% of the generated data and determines the "placement" of "structured" data to increase the incidence of frequently asked-for data being already loaded in an active tape in a library (access time =<30 seconds) or in the on-site HDD cache.
Government Public Records	Increasing, no rate specified; "Enterprise": Varies, depends on Application "Preservation": Indefinite	"Enterprise": Varibly based on the type of service, not the file size "Preservation": Divided by "small" (<1GB) and "large (>1GB)	"Enterprise" and "Preservation": ~5-7 years, depending on the platform	"Enterprise": 100% unstructured "Preservation": 100% structured
Media/ Entertainment	Strictly controlled by AI/ML archive rules and limited budget; data ingest has expanded to 8,000 hours/week, up from 2,000 hours/week in 2010	The data-object size is "all over the board," but the general "size" of each event is 250GB-300GB; the smallest files can be 1MB or smaller	Every 7 years for all architectures; year 1 is data migration/ integration, then 5-6 years of useful life, at which point the cycle begins again.	100% unstructured, updated daily; 100% of the archive is WORM, and can be actively mined for comparative content. We write on average 50TB new content into the archive per day and read 100TB-200TB from the archive for daily content production
Corporate	25%, expanding; Indefinite	Irrelevant, but we suggest a 10GB file, the larger the better for performance concerns	Usually ~5 years for all systems, no part of the infrastructure is allowed to age for much more than 7 years	100% unstructured
Large Video Surveillance Customers (with 500 or more cameras)	Depends on number of cameras and video resolution, generally >25%, but most data is deleted after 30 days of active storage;Generally <30 days, but demand/need is for at least a 90-day, preferably 6-month-to-1-year, retention time	Mostly irrelevant, depends on camera resolution	~3 years, mostly if not exclusively HDD storage systems	100% unstructured

Source: Furthur Market Research (March 2023)

Table 8: Survey Results 3

Vertical Market	% of Data Backed Up, How Many Times	Measures to Reduce Malware/Breach Risk	Measures to Reduce Unnecessary Redundancy	Measures to Reduce Power Consumption, Enhance Sustainability
Scientific 1	100% to primary on-site data center and to three other geographically dispersed global data centers	Checksum for immutability accuracy, tools to prevent any control system breach	NA, none of the scientific data gathered is considered "redundant"	Ongoing, multiple fronts, high priority
Scientific 2	100% on-prem and 100% off-prem simultaneously	Malware concerns minimal, but too-frequent access by global "bad actors" is monitored, and these actors are cut off	Provisions to prevent storing of same data from different sources, but all unique data sets are never "redundant"	Eliminated one data center, considering measures to limit on-prem backup and rely soley on off-prem for backup/archive
Scientific 3	100% on-prem; all data is essentially an "active archive " and is "updated"/"added to" with each successive run of experiments	The data is not proprietary or sensitive, so "breach" is not an issue.	NA, none of the scientific data gathered is considered "redundant," but on-site engineers constantly refine file optimization techniques	The cooling systems in the new data center more sophisticated and economical; enormous power and space savings are already achieved by using tape rather than HDDs or SSDs as the primaryTier 0 storage media.
Government Public Records	"Enterprise": 100% 1 copy "Preservation": 100% 3 copies	"Enterprise": Standard ITSEC "Preservation": Fixity Checks But both matters of public record	Largely irrelevant	Power, cooling, space constantly considered/reconsidered for all infrastructures
Media/ Entertainment	100% of "useable" (no "dead time") content to 2 archive "mirror" sites and some portion cached in production sites prior to transport to the mirror sites	Largely a non-issue	Constant use of AI/ML to analyze and prevent the unnecessary archiving of "dead time" during recorded content	Constant efforts to enhance sustainability; concentrating 7PB in a single rack reduces high-power usage to a few racks; we always look for savings in space, power, and heat generation; tape is a great saver in all of these areas and has reduced our footprint greatly
Corporate	100% daily	Breach/malware huge concerns; multiple measures constantly updated	NA, redundancy not a concern; all is kept for compliance or data mining for future business value	Ongoing, multiple fronts; current initiative is to deploy fewer servers, one big one vs several small ones, we have too many servers, by reducing server count we can vastly reduce physical space and power
Large Video Surveillance Customers (with 500 or more cameras)	No backup, all data is an "active archive" but only for ~30 days, at which point the data is deleted to provide space for new video feeds	Largely a non-issue	Largely irrelevant	Power not as great a concern as prohibitive costs to retain video feeds for more than 30 days on HDDs

Source: Furthur Market Research (March 2023)

Not surprisingly, almost all of these data center managers specified “indefinite” retention periods for the vast majority of their data, but they fear that the rising costs of preserving their data for many years or “indefinitely” will become prohibitive. (In the video surveillance arenas even the far more-modest 90-day-to-one-year desired retention span is inordinately costly.) For data retention over years or decades or centuries, not only is initial cost/GB a crucial concern, power and space and technology refresh requirements will play a far more critical role in strategic plans for future data centers.

In one large scientific data center we are aware of but did not survey, a set of new physics experiments which had to be performed on new custom architectures had been budgeted and approved, but these experiments had to be delayed and perhaps completely redesigned because the local power grid simply could not handle the additional load.

Notwithstanding the huge hype regarding massive growth in enterprise edge/hybrid cloud configurations, none of these managers indicated having any “edge/hybrid” infrastructures, but one used the off-premises cloud as a full additional backup, one used the off-premises cloud for partial backup, one used three other geographically dispersed off-premises (but non-cloud) data centers for partial backup, and one used two off-site (but non-cloud) archive mirror sites.

Somewhat unsurprisingly, because most of the data in most of these markets are matters of public record, breach/malware concerns were minimal, whereas for the corporate market breach/malware concerns were immense, and multiple defense measures have been and will be relentlessly and frequently updated.

For all these managers, data immutability was a crucial concern—all aspects of the original data absolutely must remain unchanged.

Unsurprisingly, storage requirements in all markets are increasing, even for these relatively huge databases, but the media/entertainment customer we surveyed indicated that, despite growing storage needs, budget controls and increasingly efficient archive rules strictly limit database growth. As mentioned above, the availability of power can also limit database growth.

As expected, overall 5-to-7-year infrastructure refresh cycles were normal. However, the tape robotic libraries were replaced in 10-to-15-year cycles, while the HDD and SSD infrastructures were replaced every 3-5 years.

Except for video surveillance—where HDDs are used almost exclusively, and where tape can easily decrease costs while effectively increasing retention times at higher resolution—all of these data centers use rich mixtures of SSD, HDD, and tape technology. In the case of the three large scientific community managers we interviewed, tape is widely used as a primary Tier 0 storage medium, while SSDs and HDDs are variably deployed as intelligent caches of the enormous database to decrease access times for users.

TAPE TO THE FUTURE

Unsurprisingly, none of the 15 largest hyperscale data center customers deigned to speak with us, which is unfortunate, since we estimate these customers consumed ~65% of all the annual enterprise petabytes shipped in 2022 (up from ~59% in 2020); this would equate to ~121.3K SSD petabytes, ~645.6K HDD petabytes, and ~134.2K tape petabytes delivered in 2022.

However, we know that a manager of one hyperscale infrastructure we are aware of considers incoming data to be hot for 72 hours; this data is stored on immense banks of SSDs. Once this data ages for more than 72 hours, it is migrated to far more immense banks of HDDs, where it is kept indefinitely. At the time of this recent conversation, no enterprise tape was deployed in this infrastructure.

In most of the hyperscale data centers, SSDs are used to store all of the hot and part of the warm data, and HDDs are used to store part of the warm data and virtually all of the cool, cold, and frozen data layers. A few hyperscalers have only recently become major tape customers, and accounted for most of the extraordinary 39.4% growth in tape shipments during 2021 (followed by a tepid ~9% growth in 2022, see Table 1). If current trends continue, the largest hyperscalers will continue to integrate increasingly massive quantities of SSD, HDD, and, in the future, tape technologies, and will grow to consume more than 70% of all the enterprise petabytes delivered—this could equate to these customers consuming ~8-10 zettabytes of all the SSD, HDD, and tape capacity delivered in 2030.

As previously mentioned, enterprise tape as a % of total petabytes shipped might shrink to ~11.1% in 2030 from 15.5% in 2022, despite more extensive deployments in hyperscale and other enterprise environments and a greater than 6x forecast expansion in annual petabyte deliveries 2022-2030 (from ~206.8K petabytes to ~1.3 million petabytes).

That said, because of limited production from the SSD and HDD makers, combined with compelling power and other TCO considerations—the astonishing 838x HDD + SSD vs tape 2020-2025 raw power consumption estimates detailed in Table 5 should be impossible to ignore—we believe the enterprise tape vendors will have multiple opportunities to extend tape adoption in a wider array of enterprise environments.

A prime example would be the large-scale video surveillance markets, where tape could cost-effectively provide the desired 90-day-to-one-year retention periods with acceptable access times, while HDDs cannot even cost-effectively provide 30-day retention periods, and all the HDD video data is generally erased every 30 days.

As the cold/frozen data layers expand, it will become more and more evident that HDDs and SSDs are managing far too much of the cold/frozen workloads at excessive costs/GB while sucking up a gluttonous share of available energy.

INCONCLUSIVE CONCLUSIONS

The data centers of the future will need everything the SSD, HDD, and tape industries can manufacture and deliver, as well as requiring new DNA and optical and perhaps other enterprise storage technologies, to cost-effectively and reliably preserve the priceless artifacts of our personal, corporate, and cultural history.

Availability and sustainability challenges, combined with the costs of managing our multi-millionfold-petabyte dataverse over increasingly lengthy time periods, will create new use cases for old storage technologies and demand the creation of new, more cost-effective, and power-efficient storage technologies.

Inevitably and inescapably, richly varied computing technologies will come and go, but the data we create will remain, and will grow to unimaginable immensity.

This whitepaper was jointly sponsored by Fujifilm, IBM and Twist Bioscience and was written by **John Monroe**

Furthur Market Research, LLC, **Email:** torstones@gmail.com, **Website:** Furthurdata.com
Mobile Office: +1 408-828-3967

The author wishes to express deep thanks to the IT managers who agreed to help provide some "hard data" background for this WhitePaper and deep appreciation to Brad Johns of Brad Johns Consulting LLC and Jay Bartlett of Cozaint for their guidance in tape power consumption and video surveillance markets, as well as sincere thanks and appreciation to my sponsors—Fujifilm, IBM, and Twist Bioscience.

FUJIFILM
Value from Innovation

FUJIFILM Recording Media U.S.A., Inc. delivers breakthrough data storage products based on a history of thin-film engineering and magnetic particle science such as Fujifilm's NANOCUBIC and Barium Ferrite technology. Our mission is to enable organizations to effectively manage the world's exponential data growth with innovative products and solutions, recognizing the social responsibility to protect the environment and preserve digital content for future generations.

IBM

International Business Machines Corporation is an American multinational information technology company headquartered in Armonk, New York, with operations in over 170 countries. IBM offers a full range of tape storage products including drives, autoloaders, libraries, virtual tape systems, IBM Spectrum Archive software and Hybrid solutions.

TWIST BIOSCIENCE

Twist Bioscience is a leading and rapidly growing synthetic biology and genomics company that has developed a disruptive DNA synthesis platform to industrialize the engineering of biology. The core of the platform is a proprietary technology that pioneers a new method of manufacturing synthetic DNA by "writing" DNA on a silicon chip. Twist is leveraging its unique technology to manufacture a broad range of synthetic DNA-based products, including synthetic genes, tools for next-generation sequencing (NGS) preparation, and antibody libraries for drug discovery and development. Twist is also pursuing longer-term opportunities in digital data storage in DNA and biologics drug discovery. Twist makes products for use across many industries including healthcare, industrial chemicals, agriculture and academic research.

Monroe Biography

John Monroe has been involved with the storage industry for more than 40 years, beginning in 1980.

- From October 1997 to February 2022, Monroe was a VP Analyst at Gartner. He covered the history and forecasted the future of consumer and enterprise storage markets, from components—the interplay of HDDs, SSDs, and tape—to external controller-based (ECB) networked/fabric-attached storage systems and server direct-attached storage (DAS).
- From 1990 to 1997, he was the VP of all storage lines at SYNNEX Information Technologies (now TD SYNNEX), a global distribution and manufacturing services firm, responsible for the profitable resale and OEM integration of HDDs, controllers, subsystems, and tape. Unlike most industry analysts, Monroe has had balance-sheet accountability for the stuff that he studies.
- From 1988 to 1990 he was Director of North American Sales for Kalok Corporation (a startup HDD manufacturer).
- From 1983 to 1988 he was part owner and general manager of Media Winchester, Ltd., a storage products distributor and integrator which was one of Seagate's inaugural "SuperVARs."
- He began his career in 1980 at Electrolabs, selling ICs, power supplies, cables, monitors, printers, 8-inch floppy disk drives, and 8-inch HDDs ("oddments of all things" related to computing electronics).

Monroe earned a BA degree summa cum laude, Phi Beta Kappa from Amherst College in 1976 and a master's degree in fine arts (MFA) with a merit scholarship from Columbia University in 1980.

As in his analyses and forecasts of "infinitely-self-similar-but-never-the-same" storage market trends over many years, Monroe's aim at Furthur Market Research is to bring actionable business perspectives tempered by Chaos Science, knowing that, within the unpredictably turbulent flow of dynamically changing systems—which "mirror a universe that is rough, not rounded, scabrous, not smooth," which reflect a fractal "geometry of the pitted, pocked and broken up, the twisted, tangled, and intertwined"*—there lies a deeply mysterious order that, in some way, at some scale, will always repeat itself.

**Chaos, Making a New Science*
—James Gleick

Furthur
Market Research