Furthur
Market Research

THE ESCALATING CHALLENGE OF PRESERVING ENTERPRISE DATA

Computers come and go,
but DATA remains, and it grows.

August 2022

John Monroe | Furthur Market Research

TABLE OF CONTENTS

SYNOPSIS:

Our estimated maximum production and shipment capabilities of enterprise storage vendors will fall far short of any actual 2022-2030 demand that evolves at >25%-per-year growth rates. It is becoming more and more obvious that more and more tape—as well as other forms of ultra-low-cost, massive-capacity enterprise storage technology, such as DNA data storage or new breeds of optical media—will be needed.

The preponderance of the historical evidence, coupled with deliberately conservative yet immense revenue forecasts, suggest that storage industry financial and business executives should immediately and generously fund the enhancement of old enterprise technologies and the creation of new enterprise technologies that can be deployed more cost-effectively at massive scale with minimal power consumption.

- As background research for this whitepaper, we conducted extensive interviews with three executives responsible for the management of large (>100PB) databases. Each was asked this question (among many others): ***What % of your enterprise data must be retained in some immutable form for more than 90 days? 1 year? 3 years? 5 years? 10 years? 25 years? 50 years? Indefinite?*** In each case, the answer was the same for a scientific, corporate, and media/entertainment storage infrastructure: "***Indefinite***"— for 100% of their data. In each case, the original archived data was already more than 20 years old, and was growing at escalating rates of expansion.

NOTE:

In this paper, "enterprise PB" or "enterprise data" refer to the fresh shipments and active installed base of total enterprise petabytes (PB), as defined by the specified capacity* of all enterprise-class hard-disk drives (HDDs), solid-state drives (SSDs), tape, optical, and perhaps other (TBD) building blocks of enterprise storage systems. This estimate specifically excludes PB shipments of consumer-grade HDDs, SSDs, and flash modules delivered to PCs, entertainment devices, cell phones, and other consumer and industrial applications (such as aircraft and telecom installations), the vast majority of which we believe are already backed up in, and therefore reflected by, the enterprise-grade PB serviced by corporate on-premises and cloud off-premises data centers.

*HDD capacities are raw/uncompressed, since so few enterprise HDDs utilize any form of data compression. Mission-critical (10K-rpm and 15K-rpm), performance-optimized HDDs are quickly reaching end-of-life status, and accounted for only ~1.5% of the enterprise HDD PBs delivered in 2021. Business-critical/nearline (7.2K-rpm), capacity-optimized HDDs have accounted for ~98.5% of the enterprise HDD PBs delivered in 2021 and will soon account for 100% of all HDD PBs delivered as the PC- and consumer-grade HDD markets also shrink to extinction.

SSD capacities reflect an approximate 5x compression ratio, but only for ~5% of all enterprise SSD PBs shipped, the vast majority of which (~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 is the norm.

Enterprise tape capacities are based on a 2.5x average compression ratio.

Our aim is to fairly and accurately reflect the "effective" enterprise PB under management in the installed base, above and beyond the raw enterprise PB delivered.

PROLOGUE: NOT DATA LAKES, BUT DATA OCEANS, AND DATAVERSES

The enterprise "data pools" of the early 2000s became "data lakes" by 2010 and grew in recent years to become "data oceans" which soon will morph into multiform "dataverses."

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 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.

It became obvious by 2010 (a year which might be considered to mark the true beginning of the cloud storage era) that more and more different kinds of intelligent electronic things would be connected in faster and more cost-effective networks. And it became obvious that more-diverse kinds and greater quantities of rich-media data would be exchanged in corporate and consumer systems, engendering new and enlarged dimensions of collision, convergence, dislocation and transformation. And it became blindingly obvious that most of these new dimensions in corporate and personal computing would demand more storage capacity.

Even before the "internet age," during the birth of the "PC era" in the 1980s, Al Shugart, a founder of Seagate Technology, once quipped, "The only thing I know about these damn things [PCs, and other computing devices] is that they're insatiable." For DATA, in all its digitized forms.

Increased storage at any point in the World Wide Web—bear in mind that a mobile phone is a point in the Web—increases the possibilities for storage in every part of the Web.

We are only beginning to see the enormous implications of that simple fact.

VENDOR PLANNING PERSPECTIVES

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 ~21 million freshly delivered petabytes of mostly unstructured enterprise data capacity in 2030, bringing the total active installed base of enterprise data to ~62.9 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 30.5% from 2010 through 2021 (see Table 1), 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 less than 25%. In an ultra-conservative 20%-per-annum-rate-of-increase scenario, there will be only ~7 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.

2010-2021 Historical Shipments

Table 1 below delineates historical enterprise HDD, SSD and tape shipments from 2010 through 2021. Note that while the annual petabyte deliveries grew by a 2010-2021 CAGR of 30.5%, the active installed base grew by 41.0%.

Table 1: Shipments of Enterprise HDD, SSD and Tape Petabytes

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	"CAGR 2010-2021"
Uncompressed Enterprise HDD Shipments*	45,216	52,405	65,902	90,608	116,587	157,093	217,890	261,383	333,018	485,672	679,887	998,903	
Annual Change %	–	15.9	25.8	37.5	28.7	34.7	38.7	20.0	27.4	45.8	40.0	46.9	
Compressed Enterprise SSD Shipments**	187	648	1,378	3,742	8,076	26,154	23,580	38,727	68,457	91,020	133,896	220,460	
Annual Change %	–	246.5	112.7	171.5	115.8	223.9	-9.8	64.2	76.8	33.0	47.1	64.7	
Compressed Enterprise Tape Shipments***	30,208	43,281	53,150	61,783	85,568	98,432	123,788	135,918	125,386	134,311	136,119	189,738	
Annual Change %	–	43.3	22.8	16.2	38.5	15.0	25.8	9.8	-7.7	7.1	1.3	39.4	
Total Petabytes	**75,611**	**96,334**	**120,430**	**156,133**	**210,231**	**281,679**	**365,258**	**436,028**	**526,861**	**711,003**	**949,902**	**1,409,101**	**30.5**
Annual Change %	–	27.4	25.0	29.6	34.6	34.0	29.7	19.4	20.8	35.0	33.6	48.3	
Active Installled Base	**91,000**	**180,488**	**291,129**	**430,190**	**613,881**	**819,949**	**1,088,873**	**1,404,471**	**1,775,199**	**2,275,972**	**2,944,194**	**3,988,037**	**41.0**
Annual Change %	–	98.3	61.3	47.8	42.7	33.6	32.8	29.0	26.4	28.2	29.4	35.5	
Perentage Distributions													
HDD %	59.8	54.4	54.7	58.0	55.5	55.8	59.7	59.9	63.2	68.3	71.6	70.9	
SSD %	0.2	0.7	1.1	2.4	3.8	9.3	6.5	8.9	13.0	12.8	14.1	15.6	
Tape %	40.0	44.9	44.1	39.6	40.7	34.9	33.9	31.2	23.8	18.9	14.3	13.5	

Note: The historical active installed base assumes an average five-year refresh/replacement cycle.

**HDD capacities are raw/uncompressed, since so few enterprise HDDs utilize any form of data compression.*

Mission-critical (10K-rpm and 15K-rpm), performance-optimized HDDs are quickly reaching end-of-life status, and accounted for ~1.5% of the enterprise HDD PBs delivered in 2021.

Business-critical/Nearline (7.2K-rpm), capacity-optimized enterprise HDDs have accounted for >90% of the enterprise HDD PBs delivered since 2017 and will soon account for 100% of all HDD PBs delivered, as the PC- and consumer-grade HDD markets also shrink to extinction.

***SSD capacities reflect an approximate 5x compression ratio, but only for approximately 5%-8% (the actual was 4.96% in 2021) 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 are based on an approximate 2.5x compression ratio relative to raw petabytes delivered on LTO and IBM Jaguar/ Enterprise Tape formats.*

Our aim is to fairly and accurately reflect the "effective" enterprise PB in the active installed base, above and beyond the raw enterprise PB delivered.

2022-2030 Scenarios

Devised with these precautionary adages in mind:

- *The only thing we know with certainty about any forecast is that it will be wrong. —Anonymous*
- *He who foretells the future lies, even if he tells the truth. —Arab Proverb*

The sheer amount of new scientific, corporate and consumer data being generated every day is literally incalculable. The amount of this daily explosion of unstructured data that will actually be stored and subsequently retrieved and analyzed for its potential value is also incalculable.

But we do know with rigorous accuracy the number of enterprise-grade HDD, SSD, and tape petabytes delivered historically on an annual basis. And while we cannot know with precision how many of the generated petabytes of consumer and corporate data have been and will be actually stored on enterprise-grade media and therefore cannot determine with rigorous accuracy the evolving installed base of active enterprise petabytes, we can make some hopefully accurate and intelligent guesses based on hard historical shipment data, average five-year infrastructure refresh/replacement cycles, and conservative or aggressive growth estimates.

Our beginning baseline in these scenarios of evolving demand is a compilation of actual enterprise-grade HDD, SSD and tape shipments from 2005 through 2021, an estimated active installed base of 91,000 petabytes in 2010, and an estimated five-year infrastructure refresh/replacement cycle from 2010 through 2021 (retiring all 2005 shipments in 2010, all 2010 shipments in 2015, etc., while adding successive annual shipments to the installed base of the prior year). This baseline is annually detailed in Table 1. Table 2 delineates the hard background data used in Figure 1.

Table 2: Enterprise Capacity Shipments and Active Installed Base Estimates in Exabytes and Zettabytes

	2010	2015	2021	2025	2030	"2010-2030 x Factor"
25% Shipment Growth 2022-2030, Five-Year 2022-2030 Refresh/Replacement Cycles						
Shipments	76EB	282EB	1.4ZB	3.4ZB	10.5ZB	
Active Installed Base	91EB	820EB	3.9ZB	11.5ZB	35.2ZB	387
35% Shipment Growth 2022-2030, Five-Year 2022-2030 Refresh/Replacement Cycles						
Shipments	76EB	282EB	1.4ZB	4.7ZB	21.0ZB	
Active Installed Base	91EB	820EB	3.9ZB	14.0ZB	62.9ZB	691
45% Shipment Growth 2022-2030, Five-Year 2022-2030 Refresh/Replacement Cycles						
Shipments	76EB	282EB	1.4ZB	6.2ZB	39.9ZB	
Active Installed Base	91EB	820EB	3.9ZB	16.9ZB	108.5ZB	1,193

Note: *An exabyte (EB) is a thousand petabytes and a zettabyte (ZB) is a thousand exabytes.*

Enterprise optical shipments remain minimal (to date) at ~1,500 petabytes per year—less than half of 1% of the 2021 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.

Figure 1: The Active Installed Base of Enterprise Zettabytes

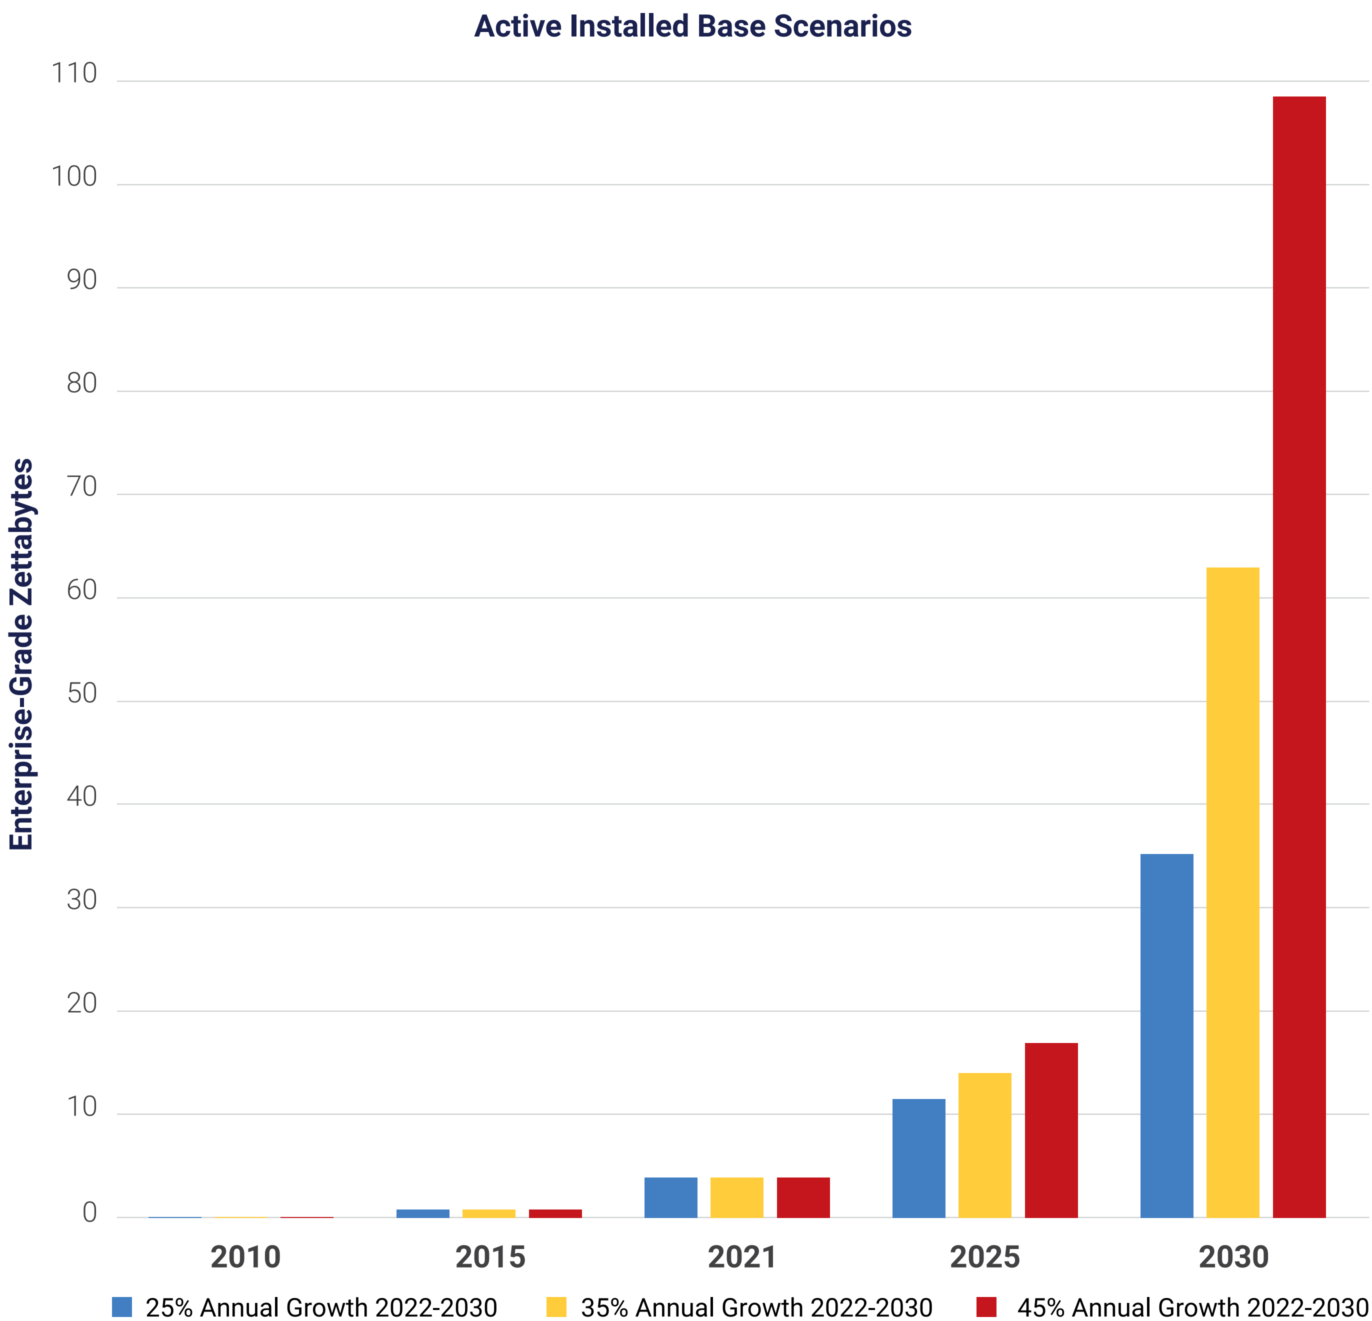

Logical assessments based on hard facts can lead to monstrous conclusions…

Note: 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.

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.

The active installed base exceeded one zettabyte in 2016 and grew to almost four zettabytes in 2021 (see Table 1). We are squarely in the midst of the zettabyte era, with the active installed base certainly exceeding 35 zettabytes but likely growing to 60 zettabytes or more in 2030, and may enter the yottabyte era in the 2030s.

For detailed 2022-2030 annual depictions of 20%, 25%, 30.5%, 35%, 40% and 45% growth scenarios, see Table 4.

In a deliberately conservative 25% annual growth scenario of potential demand (cold blue in Figure 1), the estimated active installed base grows in concert with 2010 through 2030 shipments and a five-year infrastructure refresh/replacement cycle to 35.2 zettabytes (35.2 thousand exabytes, 35.2 million petabytes) in 2030. Since shipments of enterprise petabytes are accelerating and grew respectively by 35%, 33.6%, and 48.3% in 2019, 2020, and 2021 (see Table 1), this estimate already appears to be improbably low. And yet the active installed base still grows from 2010 to 2030 by a factor of 387x.

If we assume an even more-conservative (and less-likely) 20% growth rate of potential demand from 2022-2030, we still end up with an estimated active installed base of 26 zettabytes (26 thousand exabytes, 26 million petabytes) in 2030—still a 286x expansion over 2010.

In a far more likely-case 35% growth scenario of potential demand (cautionary yellow in Figure 1), deliberately reflecting only a slight increase in the actual growth rates we saw 2010 through 2021 (30.5% over 11 years), the estimated active installed base grows to 62.9 zettabytes (62.9 thousand exabytes, 62.9 million petabytes) in 2030, an astonishing 691x expansion over 2010, a likely, if incredible, scenario.

This likely-case demand scenario might be regarded as modestly conservative, since we have seen accelerating growth rates since 2017 over an increasingly huge installed base, and there is little or no reason to believe that actual data demand will diminish in the 2022-through-2030 time period and myriad reasons to believe that data demand will continue, on average, to accelerate.

If we assume a more-conservative exact replica of the actual 30.5% 2010-2021 growth rates of potential demand from 2022-2030, we still end up with an estimated active installed base of 48.7 zettabytes (48.7 thousand exabytes, 48.7 million petabytes) in 2030—still a 535x expansion over 2010.

In a best-case aggressive but still feasible 45% growth scenario of potential demand (red danger in Figure 1)—deliberately reflecting the claim of some of the largest of the 15 global hyperscale accounts that they require 50% annual database expansion in order to achieve their strategic business aims, combined with our estimate that the 15 largest hyperscale accounts** consumed approximately 60% of all enterprise petabyte deliveries in 2021—the estimated active installed base grows to 108.5 zettabytes (108.5 thousand exabytes, 108.5 million petabytes) in 2030, a staggering 1,193x expansion over 2010. The active installed 2030 base in this scenario beggars belief, but we will surely see some years in which growth rates will exceed 40% (as we did in 2021), and a 2022-2030 45% annual growth scenario of potential demand should not be summarily dismissed.

If we assume a somewhat less-aggressive and more-feasible 40% growth scenario of potential demand, we end up with an estimated active installed base of 82.9 zettabytes (82.9 thousand exabytes, 82.9 million petabytes) in 2030, a 911x expansion over 2010.

These aggressive demand scenarios may be curtailed only by lack of available supplies of enterprise-grade media and strictly limited IT budgets. As we speculate below in Table 4, anything greater than a 25% annual 2022-2030 expansion rate in potential demand might fall into a zone of potential insufficiency.

As in all physical chaotic dynamical systems that work on themselves in iterative ways—such as galaxies or clouds or mountains or ocean currents or high- and low-pressure weather fronts—the diverse future outcomes of networked storage systems and patterns of data ingestion will evolve from variously sensitive dependence upon initial conditions, and it will always be impossible to precisely predict turbulent flows of data very far downstream in a maze of fractal basin boundaries.

There is always the danger of inferring erroneously from the possibility of concepts (logical) the possibility of things (real).

But what we can clearly see and learn from these scenarios is that logical assessments based on hard facts can lead to monstrous conclusions.

***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.*

Changing Costs, Rates of Technology Advancement, and New Forecasts Based on Production Capabilities

Table 3 details rates of ASP/GB declines for HDDs, SSDs and tape, 2010-2030. The rates of ASP/GB decline will generally slow for all three technologies as we move through the 2020-2030 decade, with HDDs and tape struggling to achieve >5% annual ASP/GB declines, and SSDs struggling to achieve >10% annual ASP/GB declines, compared with general declines in robust double-digit ranges from 2010-2021.

Table 4 details new 2022-2030 deliberately conservative annual forecasts for enterprise HDDs, SSDs and tape, based on evolving maximum production capabilities and CAPEX estimates, relative to the potential evolving demand more generally outlined in Table 2 and Figure 1 at 20%, 25%, 30.5% (the actual 2010-2021 growth rate), 35%, 40% and 45% annual petabyte expansion rates, 2022-2030.

Table 3: High-Volume Direct Ex-Factory Pricing for Enterprise HDDs, SSDs and Tape, with Available Technology

	2010	2015	2021	2025	2030
Business-Critical/Nearline HDDs					
Best ASP/GB*	$0.0590	$0.0273	$0.0108	$0.0075	$0.0061
Relative % Decline Compared to Previous Five Years	-	-53.7	-60.4	-31.8	-19.3
Average % Annual Decline Over Five Years		-10.7	-12.1	-6.4	-3.9
Available Technology	500GB/per 95mm disk	At least 1TB/ per 95mm disk	At least 2TB/ per 95mm disk	At least 3TB/ per 95mm disk	At least 5TB/ per 95mm disk
Server-Grade SSDs					
Best ASP/GB*	$1.9215	$0.3490	$0.0741	$0.0330	$0.0179
Relative % Decline Compared to Previous Five Years	-	-81.8	-78.8	-55.5	-45.8
Average % Annual Decline Over Five Years		-16.4	-15.8	-11.1	-9.2
SSD: HDD xPrice Premium	32.6	12.8	6.9	4.4	3.0
Available Technology	8 Layer BICS2	32-48 Layer BICS3	128-176 Layer BICS5	192-384 Layer BICS6?	512-768 Layer BICS8?
Enterprise Tape					
Best ASP/GB**	$0.0256	$0.0088	$0.0051	$0.0032	$0.0018
Relative % Decline Compared to Previous Five Years	-	-65.6	-42.0	-38.2	-44.4
Average % Annual Decline Over Five Years		-13.1	-8.4	-7.6	-8.9
HDD: Tape xPrice Premium	2.3	3.1	2.1	2.4	3.5
SSD: Tape xPrice Premium	75.2	39.7	14.5	10.5	10.2
Available Technology	LTO 5 3TB Compressed	LTO 7 15TB Compressed	LTO 9 45TB Compressed	LTO 11-12 180TB-360TB Compressed	LTO XX? >1PB Compressed?

**HDD and SSD ASP/GB are based on HDD and SSD vendors' ex-factory sales prices to direct global OEM customers (such as Dell and HPE) and direct large end-user customers (such as Google and AWS).*

***Tape ASP/GB estimates are based on amortized prices for tape libraries with a meaningful average number of high-capacity tape cartridges per library.*

Table 4: Capital Equipment Expenses, Maximum Production Capability of Enterprise-Grade Media, and Relative Sufficency to Potential Demand, 2010-2030

	2010	2015	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030
Enterprise HDDs												
Estimated CAPEX ($M)	$2,826	$1,718	$1,012	$1,150	$1,290	$1,210	$1,375	$1,080	$995	$970	$955	$950
Estimated Maximum Production of Enterprise Petabytes	80,870	210,550	1,285,000	1,610,000	1,990,000	2,550,000	3,150,000	3,780,000	4,609,000	5,750,000	7,015,600	8,830,000
Shipments	45,216	157,093	998,903	1,283,656	1,635,871	2,120,560	2,698,765	3,310,790	4,086,750	5,145,781	6,398,675	8,009,091
Annual % Expansion	-	34.7	46.9	28.5	27.4	29.6	27.3	22.7	23.4	25.9	24.3	25.2
Estimated Utilization %	55.9	74.6	77.7	79.7	82.2	83.2	85.7	87.6	88.7	89.5	91.2	90.7
Enterprise SSDs												
Estimated Annual NAND CAPEX ($M)	$6,685	$13,350	$27,320	$29,290	$27,908	$24,990	$28,565	$30,280	$33,739	$35,995	$31,277	$36,183
Estimated Maximum Raw Production ofTotal NAND Petabytes	11,098	92,187	638,333	812,393	1,067,127	1,441,347	1,879,623	2,359,520	2,838,804	3,551,344	4,247,407	5,014,299
Estimated Raw Petabyte Allocation to Enterprise SSD	170	21,615	189,084	223,876	295,870	401,812	548,734	777,390	1,100,661	1,491,205	2,198,336	2,935,098
Estimated Compressed Shipments ofTotal Enterprise SSD Capacity	187	26,154	220,460	268,651	355,044	482,174	658,481	932,868	1,320,793	1,789,446	2,638,003	3,522,118
Annual % Expansion of Compressed Petabytes	-	223.9	64.7	21.9	32.2	35.8	36.6	41.7	41.6	35.5	47.4	33.5
Allocation % of Raw NAND to Enterprise SSD	1.5	23.4	29.6	27.6	27.7	27.9	29.2	32.9	38.8	42.0	51.8	58.5
Enterprise Tape												
Estimated CAPEX ($M)	976	655	691	763	809	870	1,002	1,159	1,366	1,447	1,428	1,454
Compressed Shipments at 2.5x Raw Capacity	30,208	98,432	189,738	263,991	330,789	393,867	515,445	658,766	835,890	1,085,315	1,305,890	1,585,350
Annual % Expansion	-	15.0	39.4	39.1	25.3	19.1	30.9	27.8	26.9	29.8	20.3	21.4
Total Actual Historical Shipments	**75,611**	**281,679**	**1,409,101**									
Total Actual Annual Expansion %	**-**	**41.4**	**48.3**									
Percentage Distributions in the Revised Forecast												
HDD%	59.8	55.8	70.9	70.7	70.5	70.8	69.7	67.5	65.5	64.2	61.9	61.1
SSD%	0.2	9.3	15.6	14.8	15.3	16.1	17.0	19.0	21.2	22.3	25.5	26.9
Tape%	40.0	34.9	13.5	14.5	14.2	13.1	13.3	13.4	13.4	13.5	12.6	12.1
Total Estimated Potential Compressed Shipments and Revised Forecasts				**1,816,298**	**2,321,704**	**2,996,601**	**3,872,691**	**4,902,424**	**6,243,433**	**8,020,542**	**10,342,568**	**13,116,559**
Total Annual Estimated Expansion %				28.9	27.8	29.1	29.2	26.6	27.4	28.5	29.0	26.8
Potential Sufficiency/Insufficiency Relative to Demand at 20% Annual Growth 2022-2030				125,377	292,599	561,675	950,779	1,396,130	2,035,880	2,971,478	4,283,692	5,845,907
Potential Sufficiency/Insufficiency Relative to Demand at 25% Annual Growth 2022-2030				54,922	119,984	244,451	432,503	602,189	868,139	1,301,425	1,943,672	2,617,938
Potential Sufficiency/Insufficiency Relative to Demand at 30.5% Annual Growth 2022-2030				-22,579	-78,030	-135,052	-214,117	-430,860	-716,502	-1,062,173	-1,510,375	-2,351,533
Potential Sufficiency/Insufficiency Relative to Demand at 35% Annual Growth 2022-2030				-85,988	-246,383	-470,315	-807,647	-1,416,032	-2,286,482	-3,494,844	-5,203,203	-7,870,232
Potential Sufficiency/Insufficiency Relative to Demand at 40% Annual Growth 2022-2030				-156,443	-440,134	-869,972	-1,540,512	-2,676,059	-4,366,444	-6,833,285	-10,452,790	-15,996,943
Potential Sufficiency/Insufficiency Relative to Demand at 45% Annual Growth 2022-2030				-226,898	-640,931	-1,299,219	-2,356,249	-4,129,539	-6,852,913	-10,969,160	-17,192,499	-26,809,289
Total Shipments at 20% Annual Growth 2022-2030				1,690,921	2,029,105	2,434,927	2,921,912	3,506,294	4,207,553	5,049,064	6,058,876	7,270,652
Total Shipments at 25% Annual Growth 2022-2030				1,761,376	2,201,720	2,752,150	3,440,188	4,300,235	5,375,294	6,719,117	8,398,896	10,498,621
Total Shipments at 30.5% Annual Growth 2022-2030				1,838,877	2,399,734	3,131,653	4,086,807	5,333,284	6,959,935	9,082,715	11,852,944	15,468,091
Total Shipments at 35% Annual Growth 2022-2030				1,902,286	2,568,087	3,466,917	4,680,338	6,318,456	8,529,916	11,515,386	15,545,771	20,986,791
Total Shipments at 40% Annual Growth 2022-2030				1,972,741	2,761,838	3,866,573	5,413,202	7,578,483	10,609,877	14,853,827	20,795,358	29,113,502
Total Shipments at 45% Annual Growth 2022-2030				2,043,196	2,962,635	4,295,821	6,228,940	9,031,963	13,096,346	18,989,702	27,535,067	39,925,847

Note: As in Table 1, HDD capacities are raw/uncompressed, since so few enterprise HDDs utilize any form of data compression.
The HDD forecasts in this Table are deliberately conservative, with growth rates mostly in the 20% rather than 30%-to-40% ranges, while utilization of maximum production capacity increases from ~80% to ~90%.
Mission-Critical (10K-rpm and 15K-rpm), performance-optimized enterprise HDDs are quickly reaching end-of-life status, and accounted for only ~1.5% of the enterprise HDD PBs delivered in 2021. Business-Critical/Nearline (7.2K-rpm), capacity-optimized enterprise HDDs will soon account not only for 100% of all enterprise HDD PBs delivered, but also for all the HDD PBs delivered, as the PC- and consumer-grade HDD markets also soon shrink to extinction.
As in Table 1, 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.
Once a NAND fab is brought online, it must be utilized at 100%, so enterprise SSD petabyte shipments are entirely dependent on strategic allocation of the total fab capacity.
As in Table 1, tape capacities are based on an approximate 2.5x compression ratio relative to raw petabytes delivered on LTO and IBM Jaguar/Enterprise Tape formats.
Tape forecasts in this table assume increasing market acceptance and consistent technology execution resulting in maximum production and shipment of compressed enterprise tape capacity growing to more than 1.5 zettabytes in 2030.

To date, available enterprise storage capacity has not been a bottleneck; seemingly, the petabytes are always available, increasingly cost-effective, and seldom utilized to full capacity. But this may not be the case in the palpably near future.

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 2022-2030 than in the 2010-2021 time period, and maximum available capacities may be limited. This may engender disruptive shortages in the sum total of available enterprise petabytes relative to a real evolving demand.

Even if enterprise HDD, SSD and tape shipment capabilities grow to ~8 million (uncompressed), ~3.5 million (compressed), and ~1.6 million (compressed) petabytes per year, respectively, these advancing 2022-2030 estimated production capabilities will fall far short of any actual 2022-2030 demand that evolves at greater than 25%-per-year growth rates. As indicated above, we believe an overall 25% per year growth rate is "improbably low," but in these revised conservative forecasts based on maximum production estimates we have deliberately kept total 2025-2030 expansion rates of estimated shipments in the 20% rather than 30% ranges.

In our most-likely 35% demand growth scenario, with new shipments of 4.7 million petabytes in 2025 and 21 million petabytes in 2030, the industry in this revised forecast falls behind potential demand by ~85,988 petabytes in 2022 (this year). By 2030, the zone of potential insufficiency grows to 7.9 million petabytes. Our aggressive (but still feasible) 40% and 45% growth scenarios project zones of potential insufficiency growing to more than 15 million petabytes.

If 2022-2030 annual growth rates merely mimic the actual 2010-2021 growth rates at 30.5%, with new shipments of 4.1 million petabytes in 2025 and 15.5 million petabytes in 2030, the industry in this revised forecast falls behind potential demand by ~22,579 petabytes in 2022 (this year). By 2030, the zone of potential insufficiency might grow to 2.4 million petabytes.

In all scenarios, we believe most of this enterprise data will be unstructured, "cold," infrequently accessed, and will have to be maintained at minimal cost (see Figure 2 and Figure 3).

It is becoming more and more obvious that more and more tape—as well as other forms of ultra-low-cost, massive-capacity enterprise storage technology, such as DNA data storage or new breeds of optical devices—will be needed.

Enduring Questions, Mercurial Answers

Our 2022-2030 scenarios engender critical questions—which must be asked and answered again and again by storage industry engineers and business executives—regarding the potential evolution of enterprise HDD, SSD and tape technology, production and pricing metrics.

Perhaps the most important question—in conjunction with evolving maximum production rates—is the final one: ***What percent of limited IT budgets will be allocated to storage technologies?***

- ***Will the HDD industry be able to drive ASP/GB down to $0.006 or less?*** The best ASP/GB for a business-critical/nearline HDD during 4Q21 was $0.0108. While HDD technologists have doubts about profitably driving ASPs/GB below $0.005, ~$0.006 ASPs/GB figure in their 2025 through 2030 strategic game plans. We regard 2025-2030 ~$0.0075 ASPs/GB to be a virtual certainty as per-drive capacities advance from 20TB to 50TB.
- ***Will the HDD industry be able to manufacture more than 8 million enterprise petabytes (8,000 exabytes, 8 zettabytes) in 2030 (a ~5x increase over current maximum production estimates)?*** Perhaps even 5 or 6 million petabytes will not be realistic, given the HDD makers' fiscal concerns—tempered by a colorful history of profitless price wars caused by needless surplus production that rivaled the regularity of the seasons—about investing unprofitably in future CAPEX in the face of uncertain demand and growing SSD incursions. In addition to funding expensive head and media manufacturing facilities, another mounting cost will be the purchase of adequate numbers of testing "beds" to expedite timely product delivery. It takes 4 days to burn in a 1TB HDD after it exits the clean room preparatory to packing and shipment. Irrespective of PC- or enterprise-grade specifications, the math is the same as capacities increase. It now takes 80 days to burn in a 20TB business-critical/nearline HDD and will require—absent some technology breakthrough, or customer-approved relaxed reliability standards—200 days to burn in 50TB drives. ***The key takeaway here being: Although the HDD forecast in Table 4 is deliberately conservative, with annual growth rates mostly in the 20% rather than 30% ranges, it still may be unrealistic relative to maximum production capacities. Since we forecast HDDs to account >60% of the available capacity, this means that the 2025-2030 zones of potential insufficiency relative to demand could be even greater.***
- ***Can SSD ASPs/GB decline at a faster pace during the 2020s?*** Our current estimates foresee a 3:1 best available ASP/GB differential—$0.0179:$0.0061—in favor of enterprise HDDs in 2030. SSDs will not reach this 3:1 cost/GB differential—the presumed ratio, as yet unproved, at which there might be a tectonic shift in favor of SSDs—prior to 2030. Even if HDD ASPs/GB remain stagnant at, say, $0.0075 2025-2030 (an unlikely scenario), there will still be a 4.4:1 differential in favor of HDDs in 2025 and a 2.4:1 differential in favor of enterprise HDDs in 2030.
- ***Despite planned fab expansion in China and elsewhere, will the NAND industry actually build the fabs necessary to deliver more than 5 million total petabytes in 2030?*** Post-COVID supply constraints, coupled with inconsistent demand and recent chaotic price metrics, have delayed NAND capacity additions and temporarily idled new equipment ramps (NAND CAPEX is now forecast to decline in 2023 and 2024 before recovering in 2025 and mostly ramping 2026-2030). Future cycles of supply/demand imbalances and the attendant NAND resale price increases and erosions will inevitably impact planned CAPEX spending and the timing of bringing new fab capacity online.
- ***Can the NAND industry afford to allocate more than 50% of its total NAND production solely to enterprise SSDs (in addition to satisfying NAND demand from the mobile phone, PC and other markets)?*** The enterprise SSD markets are the most lucrative for the NAND vendors, and they should receive favorable allocation during times of tight supply, but multiple end markets must still be supported, whereas the HDD makers will soon have only one end market (enterprise data) to support, and one class of product (business-critical/nearline HDDs).
- ***When will the NAND industry be able to deliver 5 million or more petabytes allocated solely to enterprise SSDs? At what cost? How much additional CAPEX will be required to enable enterprise SSDs to account for more than 30% of the total enterprise petabytes delivered in 2030?*** Total NAND 2022 through 2030 CAPEX is expected to

exceed $275 billion (see Table 4). According to our current forecasts, this will enable enterprise SSDs to account for 26.9% of the enterprise petabytes delivered in 2030, up from 15.6% in 2021.

- *SSD executives often demand of their engineers and salesfolk (and industry analysts): When will we be able to entirely displace enterprise HDDs?* The likely answer is not before 2035 (if then), but we believe the question itself is wrongheaded, since any meaningful estimate of enterprise needs will involve not "either/or" but rather "both/and/and" scenarios of technology consumption. It will not be a choice between, but of, available technologies configured in concert and increasingly seamless conjunction with each other—and if the needed capacity of any technology is not available at whatever cost, one cannot integrate it. What the SSD executives should be asking is: How many of the evolving millions of enterprise petabytes that will surely need to be serviced can we profitably address? When?
- *Will the tape industry achieve its technology transitions more predictably and consistently 2022-2030 and be able to drive ASPs/GB down to less than $0.002 and expand its production capacity to 1.5 million or more petabytes (1,500 exabytes, 1.5 zettabytes) per year?* The aim of the tape suppliers is to maintain some "order of magnitude" difference between HDD and tape pricing, and the technology seems feasible and affordable from a development perspective. 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 20TB uncompressed for enterprise HDDs, and should certainly deliver >300TB compressed per cartridge prior to 2030, whereas 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 new breeds of tape or other inexpensive enterprise storage technologies should play a crucial role here, potentially servicing millions of those evolving petabytes of demand (especially in the cold-yet-accessible archive arenas) and taking some pressure off required HDD and SSD capacities.
- *Irrespective of petabytes delivered on any media platform, how much enterprise storage will actually be needed and consumed? How can it be most effectively configured and managed? Can growth rates be curtailed? Is there any way to put the storage beast on a diet? What percent of limited IT budgets will be allocated to storage technologies?* If data is in fact the new "oil" and is valued as such, then the portion of IT budgets allocated by CFOs to storage technologies must grow to meet expanding enterprise needs, especially if supply shortages cause price increases. If the budget is not available, then CIOs must learn to hit the delete key without causing any deleterious impact to their companies' competitive fitness.*** To date, we have seen little or no inclination of any kind to hit the delete key, and there is scant reason to believe that this trend will change. In fact, for most organizations, the length of data retention in some immutable form seems to be "indefinite."****

****Several enterprise IT managers with whom we spoke stated that an exacerbated problem is establishing generally agreed-upon ground rules for deleting data. When they asked for buy-in from their internal clients, they could not obtain solid commitment for 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.*

******As background research for this whitepaper, we conducted extensive interviews with three executives responsible for the management of large (>100PB) databases. Each was asked this question (among many others): ***What % of your enterprise data must be retained in some immutable form for more than 90 days? 1 year? 3 years? 5 years? 10 years? 25 years? 50 years? Indefinite?*** In each case, the answer was the same for a scientific, corporate, and media/entertainment storage infrastructure: "***Indefinite***"— for 100% of their data. In each case, the original archived data was already more than 20 years old, and was growing at escalating rates of expansion.*

The Evolving Enterprise Storage Pyramid, Relative "Temperature" of the Data Layers, and the Birth of the Deep-but-Accessible Archive

Although the recently coined term "active archive" has all the ingredients for maximum hype and minimal hard data, we believe the term points to a crucial trend. It is obvious that the mission-critical, "hot" data layers of the storage pyramid depict a competitive need for "business at the speed of flash." In the social-networking world and increasingly in the traditional business world, the longer it takes to tap into a bit of data, the less valuable that data becomes. And the value of any data seems to diminish greatly after about a five-second wait time—we have become a "spectating" global culture with an egregious lack of patience.

And yet, because digital data in all its myriad forms has become the new "oil," mining and deriving business value at acceptable (rather than lightning-quick) speeds and minimal costs from the "archive avalanche" of cold, unstructured data is becoming an ever more-critical competitive need.

The so-called "storage pyramid" has been around since the 1960s and has reflected diverse and changing assumptions over the years. Initially, tape comprised the largest "cold" segments at the middle and bottom of the pyramid, and emerging HDD technology accounted for almost all of the "hot" performance-intensive segments at the narrow peak of the pyramid, with a minuscule segment for solid-state storage at the apex of the pyramid.

If we divide the pyramid today, the segments would be vastly different. Based on actual 2021 petabyte shipments (see Table 1), SRAM, DRAM and storage-class memory (SCM) technologies would comprise about the same minuscule <1% segment at the apex of the pyramid, but NAND-flash-based enterprise SSDs would comprise a growing ~15.6% segment and enterprise HDDs the majority ~70.9% segment, while enterprise tape would comprise a ~13.5% segment.

Tape's share of the enterprise petabyte total has declined drastically in recent years, to ~13.5% in 2021 from ~34.9% in 2015. This is largely explained by the extraordinary growth in hyperscale consumption, wherein HDDs have been used as the primary nearline and backup devices, and is now changing. Recently, hyperscalers have become the largest consumers of enterprise tape. As implied above, we believe there will be a resurgence in tape shipments for a variety of reasons based on expanding demand in multiple markets (as we saw in 2021, after years of tepid growth and actual decline 2017-2020, see Table 1), relative data temperature and time-to-data needs, and lower costs of data retention and power consumption, as well as limited HDD and SSD production capabilities.

Figure 2 depicts an updated version of the storage pyramid, with an ever-widening base of unstructured data. It will be increasingly difficult to imagine and depict the widening base of the deep-but-accessible archive.

Figure 2: The Evolving Enterprise Storage Pyramid

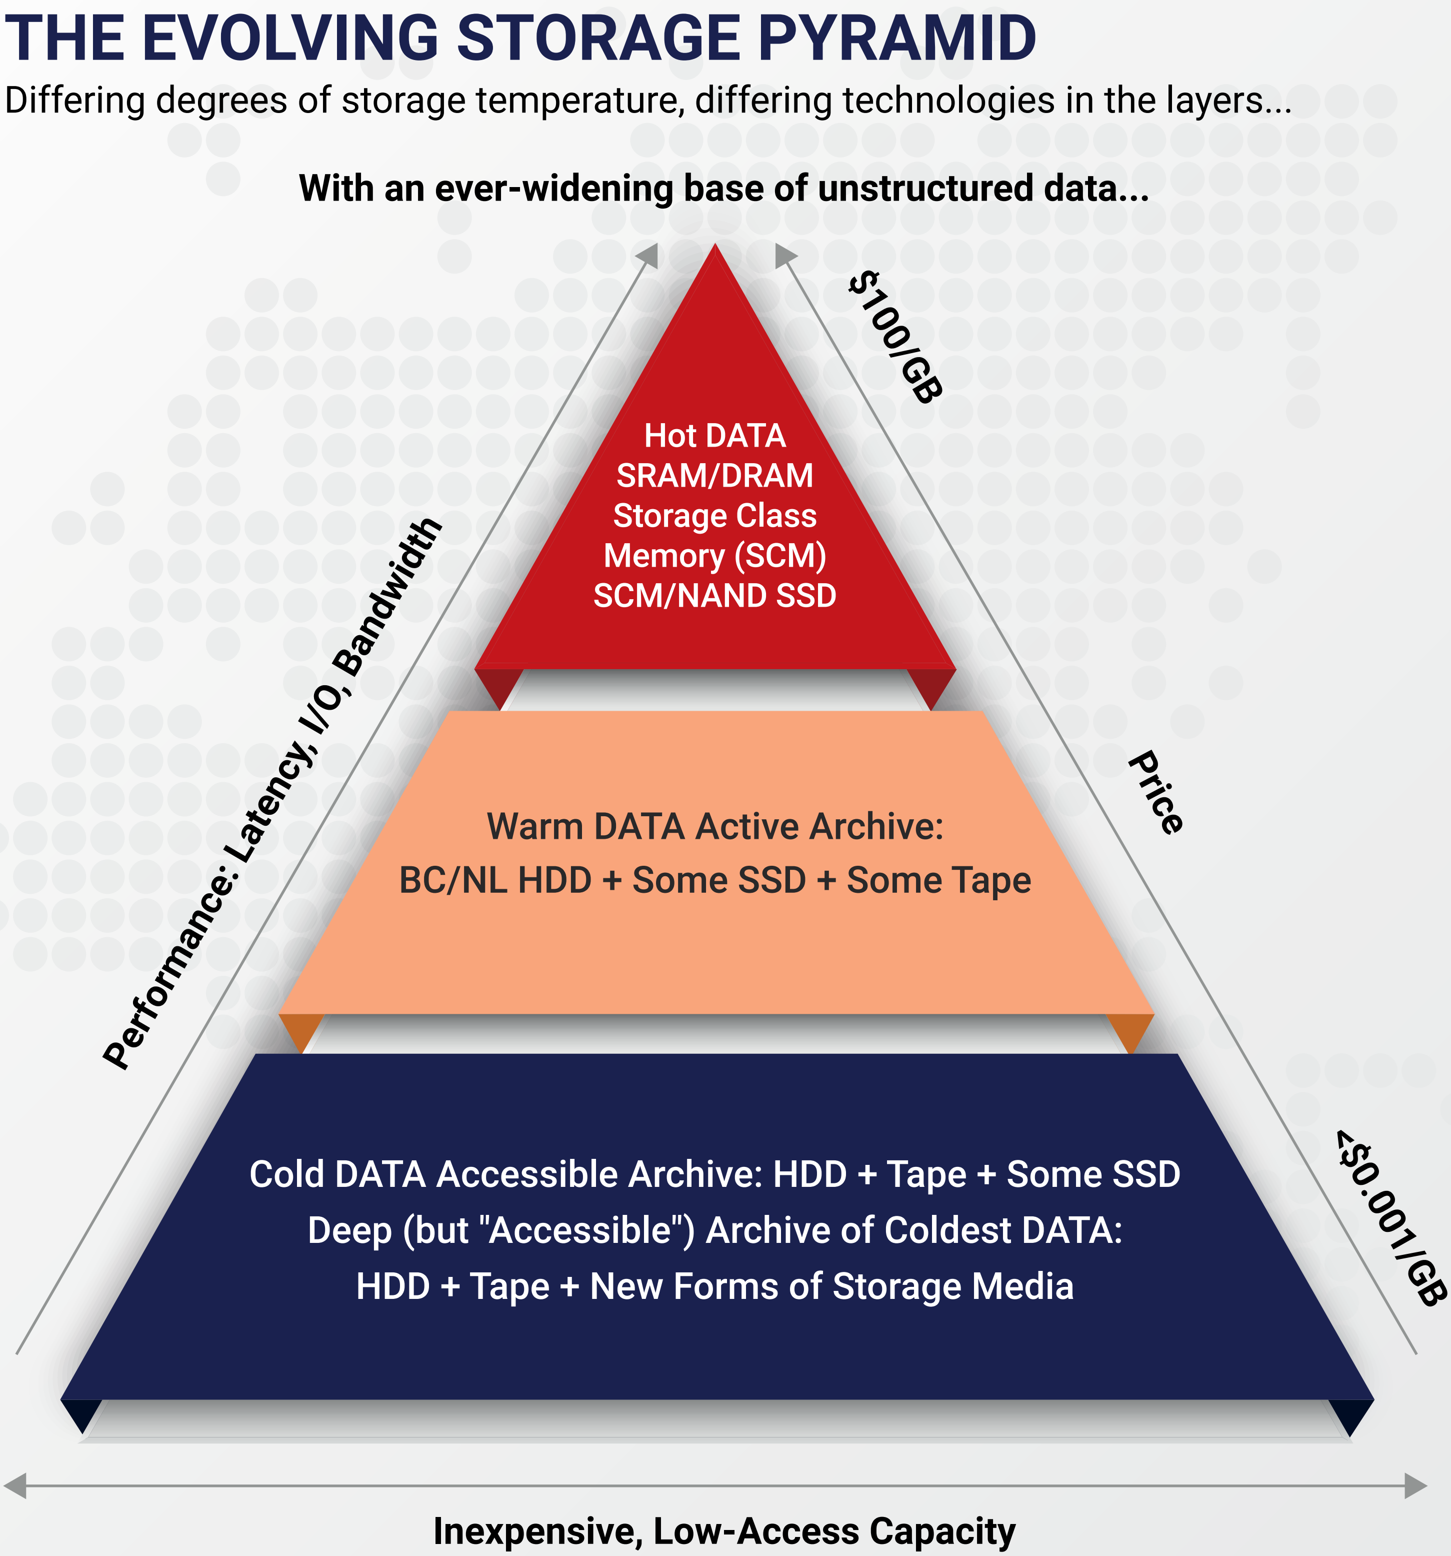

We can generally (and perhaps erroneously) regard the top ~8% of all this enterprise data as "hot" (incessantly accessed at nano-second-to-micro-second times to data), with ~17% as "warm" (frequently accessed at micro-second-to-milli-second times to data), and ~75% as "cold" (infrequently accessed with milli-second-to-five-second-to-48+-hour times to data). A majority of the "cold" data may in fact never be accessed—nor in most cases will it ever be deleted. Also, we can generally (and perhaps erroneously) regard ~28% of this data ("hot" plus "warm" plus a small portion of the "cold" layer) as "structured" and ~72% (all "cold") as "unstructured."

These are of course vast oversimplifications and may not truly reflect changing requirements in the data layers for many users. For example, one large storage infrastructure that we are aware of uses tape as the direct primary rather than secondary or tertiary medium for storing massive amounts of generated data, which is subsequently tapped into by users equipped with faster HDD or SSD technologies. In any case, the largest customers are now demanding (however unreasonably) limitless capacity and ASPs declining to less than $0.001/GB for their unstructured, deep-but-accessible archives.

Perhaps an easier and more compelling way to visualize our dataverse is as a multi-dimensional fractal iceberg—enmattering trillions of write/read/restore iterations in its trillion-fold chaotic depths and surfaces—rather than as a simple two-dimensional pyramid, with the majority of its mass lying huge and invisible beneath the water surface, constantly shaped and re-shaped by data exchanges in the surrounding oceanic global networks.

Figure 3: Our Fractal Dataverse

Let us assume that ~75% of the active installed base will soon be (if not already) "cold." If this estimate is accurate, and if our most-likely 35% annual growth estimates for 2022-2030 shipments prove to be true, then the "cold" nearline backup and archive layers will comprise ~4.1 zettabytes in 2022 and will expand to ~10.5 zettabytes in 2025 and to ~47.1 zettabytes in 2030.

Vendor Revenue Opportunities

Enterprise storage has been dismissed by some senior industry analysts as stagnant, predictable, and boring.

While it is true that, ironically, the most dynamic segments of the evolving enterprise storage markets will involve the escalating challenges of managing mostly static data, it is also true that symbiotic and lucrative business opportunities for enterprise SSDs, HDDs, and tape—as well as for DNA data storage and optical and surely other yet-to-be-conceived technologies, not to mention storage-as-a-service (STaaS) opportunities—will be immense. Table 5 attempts to conservatively delineate some of these business opportunities.

Table 5: Conservative Potential Vendor Enterprise Storage Revenue Opportunities Based on Best (Not Average) ASPs/GB, 2022-2030

	2022	2023	2024	2025	2026	2027	2028	2029	2030
Enterprise SSD Revenue ($B)	$19.6	$18.9	$18.4	$22.0	$32.0	$38.2	$35.1	$47.0	$63.0
Compressed SSD Petabyte Delivered*	268,651	355,044	482,174	658,481	932,868	1,320,793	1,789,446	2,638,003	3,522,118
Best SSD ASP/GB ($)	$0.0731	$0.0532	$0.0382	$0.0334	$0.0343	$0.0289	$0.0196	$0.0178	$0.0179
Enterprise HDD Revenue ($B)	$12.8	$14.2	$16.8	$20.2	$22.8	$27.4	$33.4	$40.3	$48.9
Raw HDD Petabytes Delivered*	1,283,656	1,635,871	2,120,560	2,698,765	3,310,790	4,086,750	5,145,781	6,398,675	8,009,091
Best HDD ASP/GB ($)	$0.0100	$0.0087	$0.0079	$0.0075	$0.0069	$0.0067	$0.0065	$0.0063	$0.0061
Enterprise Tape Revenue ($B)	$1.8	$2.0	$2.1	$1.6	$1.9	$2.2	$2.5	$2.7	$2.9
Compressed Tape Petabytes Delivered*	263,991	330,789	393,867	515,445	658,766	835,890	1,085,315	1,305,890	1,585,350
Best Tape ASP/GB ($)	$0.0068	$0.0061	$0.0053	$0.0032	$0.0029	$0.0026	$0.0023	$0.0021	$0.0018
Other Technology Revenue ($B)	$0.3	$0.8	$1.4	$1.9	$2.7	$3.7	$4.5	$5.7	$7.1
Other Technology Petabyte Demand**	85,988	246,383	470,315	807,647	1,416,032	2,286,482	3,494,844	5,203,203	7,870,232
Best Other Technology ASP/GB ($)	$0.0035	$0.0032	$0.0029	$0.0023	$0.0019	$0.0016	$0.0013	$0.0011	$0.0009
Total Revenue ($B)	$34.6	$35.9	$38.6	$45.7	$59.4	$71.4	$75.6	$95.7	$121.8
Potential Total Petabyte Demand***	1,902,286	2,568,087	3,466,917	4,680,338	6,318,456	8,529,916	11,515,386	15,545,771	20,986,791

**SSD, HDD and Tape segments directly reflect forecast shipment estimates in Table 4.*

***The Other Technology segment represents the zone of potential insuffiency delineated in Table 4 that might be serviced by enhancements of old techologies or deployments of new technologies if actual 2022-2030 demand grows at 35% per year.*

Tape product planners and marketing executives should consider some significant portion of this "other technology" segment as potential arenas for new breeds of enterprise tape.

****Total petabytes delivered include the Other Technology segment, which is added to the Table 4 HDD, SSD and Tape shipment estimates to reflect a 35% annual growth estimate.*

Discounting the significant but proportionally small revenue generated by expensive enterprise DRAM, SRAM, and SCM technologies, consider the revenue opportunities for cost-effective enterprise SSDs, HDDs, and tape based on shipment estimates in Table 4. Also, consider the upside revenue opportunities for new tape or other technologies that might service at ~$.0035-$0.0009 ASPs/GB—to repeat, the prices now being demanded, however unreasonably, by the largest enterprise customers, which may not be achievable with HDD or SSD technologies before 2040 (if then)—those "zones of potential insufficiency" depicted by the red petabytes in Table 4 in a 35% growth scenario.

Industry average rather than best ASPs/GB—available only to the largest enterprise customers—should increase these revenue opportunities by at least 10% to include specialized components or higher performance drive options; for example, enterprise revenue for the SSD makers in 2025 will likely exceed $30B (up from ~$22B in 2021). And these revenue estimates should increase by at least 25% if we consider markups for integrated storage systems rather than merely for the essential storage building blocks.

But even at these aggressive, "unreasonable" ASPs/GB, vendor revenue could grow from $35B in 2022 to more than $120B in 2030. And STaaS revenue—on-premises or off-premises consumption-based pricing services—generated from the petabtyes provided by these essential building blocks of enterprise storage may grow to >$50B in 2030 on top of these estimates, up from ~$8B*** in 2021.**

*****We believe the cloud-data-service revenue generated by just the largest hyperscalers exceeded $5B in 2021.*

Table 6 attempts to size only the "coldest" backup/archive markets, which might comprise 55% of total enterprise petabyte shipments, but must be serviced at ultra-low ASPs/GB. Even at these ASPs/GB—which, while slightly higher than the Other Technology ASPs/GB in Table 5, may prove to be profitless for any vendor—backup/archive revenue opportunities could grow from $6.8B in 2022 to $17.3B in 2030.

Note that if we, say, double the backup/archive ASPs/GB and halve the addressable petabytes, the projected revenue remains the same. Many revenue scenarios are possible, but all foretell potentially profitable growth.

Table 6: Potential Vendor Enterprise Backup/Archive Revenue Opportunities Derived from Servicing Only the Coldest Layers of Untructured Data, 2022-2030

	2022	2023	2024	2025	2026	2027	2028	2029	2030
Coldest Backup/Archive Revenue ($B)	$6.8	$7.2	$7.8	$9.3	$10.1	$11.7	$13.3	$15.4	$17.3
Aggressive ASPs/GB ($)	$0.0065	$0.0051	$0.0041	$0.0036	$0.0029	$0.0025	$0.0021	$0.0018	$0.0015
Coldest Backup/Archive Layers, 55% of Total	1,046,257	1,412,448	1,906,804	2,574,186	3,475,151	4,691,454	6,333,462	8,550,174	11,542,735
Total Petabytes Delivered*	1,902,286	2,568,087	3,466,917	4,680,338	6,318,456	8,529,916	11,515,386	15,545,771	20,986,791

**Based on 35%-per-year growth estimates in Table 4 and Table 5.*

The phenomenal revenue estimates in Table 5 and Table 6 demand that we ask again a crucial question we posed earlier, namely: ***What percent of limited IT budgets will be allocated to storage technologies?***

Irrespective of budget constraints relative to actual demand and the limits of technology advancements and maximum production capabilities, the preponderance of the historical evidence, coupled with deliberately conservative revenue forecasts, suggest that storage industry financial and business executives should immediately and generously fund the enhancement of old enterprise technologies and the creation of new enterprise technologies that can be deployed more cost-effectively at massive scale with minimal power consumption.

END USER PLANNING PERSPECTIVES

Questions of Power and Long-Term Costs to Store, Maintain, and Service Enterprise Data

Obviously, new storage infrastructure technologies must consume less power and allow for scalably efficient energy consumption. As with the production of enterprise-grade media, the need to integrate storage technologies in new ways will be crucial to be in alignment with the total availability of energy. Moreover, clients are now analyzing storage vendors with greater scrutiny and making purchasing decisions based on the sustainable energy efficiency of a vendor's products.

As the world's appetite for data increases, the challenges will grow correspondingly. Recent research by the Uptime Institute has shown that data center power usage effectiveness (PUE) made great progress from 2007 through 2018, improving from 2.51 to 1.58, but has since stalled and, in fact, increased to 1.67 in 2019. (Lower PUE indices are better.) Because of the rapid expansion of digital data needs, storage as a percentage of data center energy consumption will continue to expand and could account for more than 35% of total data center power requirements in 2030, up from ~18% in 2020.

While HDD and SSD power efficiency per-terabyte increases marginally as per-drive capacities increase, the power consumption per drive remains fairly constant at an average of ~5.4/5.8 idle watts (SATA/SAS) for business-critical/nearline high-capacity HDDs and ~1.5 idle watts (SATA) for high-capacity enterprise SSDs (<10% of enterprise SSDs have SAS interfaces). In fact, this power consumption per drive will increase for HDDs as we move to higher capacities, more disks per drive, and dual-actuators, and for SSDs with the increasing advent of nonvolatile memory express (NVMe) interfaces, with an average power draw of 3.5 watts in idle mode. (NVMe interfaces might also be integrated in significant numbers of future HDDs.) Active rather than idle power consumption for business-critical/nearline high-capacity HDDs is ~6.4/9.4 read/write watts per drive for SATA interfaces and ~7.0/9.8 read/write watts per drive for SAS interfaces. Active rather than idle power consumption for high-capacity enterprise SSDs is ~2.1/3.2 read/write watts per drive for SATA interfaces and ~11.0/13.5 read/write watts per drive for NVMe interfaces, and we estimate that NVMe interfaces will account for ~80% of all near-term enterprise SSD shipments.

Assuming 65%/35% idle/active use patterns—which may be an aggressive estimate, HDDs in the warm and cold data layers may be utilized at less than 25% in active mode—for ~408 million SATA-interface drives and ~72 million SAS-interface drives that will be shipped and integrated 2020-2025, we estimate that these enterprise HDDs will likely consume a minimum of ~2,417 megawatts of additional power (using read not write active watt specifications). Assuming 40%/60% idle/active use patterns—which may be a conservative estimate, SSDs in the hot data layers may be utilized at more than 70% in active mode—for ~90 million SATA/SAS-interface drives and 360 million NVMe-interface drives that will be shipped and integrated 2020-2025, we estimate that these enterprise SSDs will likely consume a minimum of ~3,047 megawatts of additional power (using read not write active watt specifications).

Regardless of substantial rack and cooling and other infrastructure costs, just powering these essential storage building blocks will require at least ~5,464 megawatts, equivalent to the annual electricity needs of ~438 million homes. 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, at least 16,392 megawatts of power will be consumed 2020 through 2025 to fuel the online activities of enterprise HDDs and SSDs.

Tape to the Future

As previously stated, 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, and lower costs of data retention and power consumption, as well as limited HDD and SSD production capabilities.

Enterprise tape not only will have lower initial acquisition costs/GB than HDDs or SSDs 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 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), enterprise tape will likely consume less than 18 megawatts of additional power from 2020 through 2025 and, unlike HDDs and SSDs, will be largely impervious to malware attacks, because of "offline air-gap" protection of sensitive enterprise data. Also unlike enterprise HDDs and SSDs, enterprise tape is easily removable and transportable and is delivered with standard data encryption and can function as an indelible write-once, read-many (WORM) device, to satisfy a broad range of medical and legal compliance requirements. Encryption (but not WORM) is available for enterprise HDDs and SSDs, but it is costly, and many customers do not want to pay the premium.

The cost and availability of power and space to sustain enterprise data requirements, and the long-term reliability and durability of the storage media, will become more-crucial concerns during the 2020s. The immutability of the data will also become a more-crucial concern.

Currently, enterprise tape has a specified shelf life of ~50 years—assuming storage conditions of 61 to 77 degrees Fahrenheit (16 to 25 degrees Celsius), and 20% to 50% relative humidity. Though boasting specifications of millions of hours mean time between failures (MTBF), enterprise HDDs and SSDs are designed for 100% online duty cycles with limited encryption and immutability attributes; are seldom deployed for more than five years and often for only three years, which means the data must be continuously re-mastered; and do not fare well powered down and sitting on shelves, and therefore offer no "offline air-gap" protection from malware attacks. Tape will also be plagued with remastering requirements, but they should occur over a longer time span (at least five to seven years) at significantly less cost than for HDDs and SSDs.

It should be noted that HDDs and SSDs—because of the design of the industry-standard SATA, SAS and NVMe interfaces—will have minimal backward compatibility issues as technologies advance, while tape has increasing backward compatibility issues as technologies advance. For example, LTO was intended to be backward-compatible for two generations of read and one generation of write, but LTO-9 will only read and write LTO-8 and LTO-9 data formats.

Tape's greatest comparative drawback is, of course, access time to data, which is improving but is still much slower than for HDDs and SSDs—seconds (if the tape is in the drive, minutes to hours or longer if not), as opposed to milli- to micro- to nano-seconds. This can be partially mitigated by integrating tape in close concert with SSDs and HDDs, either in the same storage system, or nearby, connected by some kind of high-speed data fabric.

Emerging Enterprise Storage Alternatives: Optical, Holography, and DNA Data Storage

Optical Technology Is the Technology of the Future, and It Always Will Be (or Maybe Not)

Another candidate for filling the potentially huge backup/archive gaps is enterprise optical, which has been around since the 1990s but has yet to gain any strategic traction, largely due to high costs and the useless and fierce contentions of various camps of incompatible technologies.

Optical discs offer a rugged, reliable, removable, transportable, random-access medium for enterprise storage and can (theoretically) last for hundreds of years (much longer than tape) and withstand (theoretically) extreme degrees of temperature and environment, from freezing to fire to boiling in water, with much greater resistance to dust and humidity than tape. Optical technologies should have equivalent or better long-term retention costs than tape and offer similar immutability and "offline air-gap" data-protection attributes, but relative initial acquisition costs are and may remain—because of uncompetitive economies of scale—inordinately high compared with tape (or HDDs, or even SSDs).

Sony, in partnership with Panasonic, began shipping Archival Disc (AD)—an extensive reengineering of its prior ProDATA enterprise technology—in 2015 with an initial capacity of 300 gigabytes per cartridge (with 11 27-gigabyte discs per cartridge). Deploying three optical layers per side per disc, AD in recent years has advanced from 1.5 to 3.3 to 5.5 terabytes per cartridge (with 11 500-gigabyte discs per cartridge). Although reputedly working on 11-terabyte cartridges, the delivery date of any AD capacities beyond currently shipping 5.5-terabyte cartridges remains unclear, and plans to reduce future integration costs may prove to be untenable. Some of the strategic, less cost-conscious customers are notable (for example, a few major Hollywood studios and the Vatican), but shipments are declining, and AD is compatible only with itself.

An alternate approach that may hold great future promise is Microsoft's Project Silica, funded in parallel with Microsoft's DNA data storage research (it should be noted that Microsoft seems to be paying much more attention to and allocating more funds to DNA data storage). Microsoft is using ultrafast femtosecond lasers that are commonly used in LASIK surgery to permanently change the structure of quartz glass by creating, at various depths and angles, layers of three-dimensional nanoscale gratings and deformations. When polarized light shines through the glass, complexly illumined figures appear that machine learning/artificial intelligence (ML/AI) firmware decodes and reads back as digital data. It should be noted that, as with holography, precisely locating any bit of data that has been written in a three-dimensional (as opposed to multi-layer) geometry is dismayingly difficult and time-consuming. To date, only the writing of the original 1978 Superman movie onto glass—an unimpressive (in fact, minuscule) amount of data—has been put forth publicly as proof of concept, and the technology, like AD, is compatible only with itself. We do not expect cost-effective commercial viability of Microsoft's Project Silica technology to occur before 2028.

Folio Photonics (FP) is a small company that was founded in 2012 as a spin-off from the Case Western Reserve University Center for Layered Polymeric Systems. FP now seems poised to deliver its new optical technology, enabling eight or 16 film layers per side per disc (as opposed to only three optical layers per side per disc for AD). Using next-generation materials, patented polymer extrusion, and film-based disc construction processes (distinct from mere optical layering), in concert with customized optical pickup units (OPUs), FP is (theoretically) engendering far greater data densities than was thought possible several years ago. The company plans to ship sample 10-to-20 terabyte cartridges (containing 10 1-to-2 terabyte discs) by 2024. Unlike AD and Project Silica, FP's technology is designed to heavily leverage existing Blu-ray disc manufacturing infrastructures and be backward read compatible with Blu-ray disc (and perhaps even older DVD and CD formats), and has a published roadmap extending to 50 terabyte cartridges by 2028 and 100 terabyte cartridges by 2033, with complete backward compatibility. We do not expect cost-effective commercial viability of FP's technology to occur before 2025.

Holography

Holographic storage, in which writing and reference laser beams store data in densely latticed patterns scattered throughout an optically sensitive medium, held hope in the early and mid-1990s as an innovative, lightning-fast, rugged technology that contained no moving mechanical parts and could generate either fixed or removable media of enormous (as yet undetermined) storage capacity. Several well-funded consortia explored this technology (for example, InPhase was backed for a while by Hewlett Packard Enterprise), but a working, cost-effective, manufacturable prototype was never born. Because the idea is so attractive, holography is still much discussed at conferences and elsewhere, but funding and research efforts have ceased (at least temporarily), and we do not believe holography will be a market force in the foreseeable future.

DNA Data Storage: "Our DNA is, after all, who we are…"

Unlike optical, DNA data storage has generated not only enormous interest but also immense funding in recent years. Also unlike optical, many global companies have chosen to participate in the synergetic development and dissemination of DNA data storage technologies. There are now more than 50 members of the DNA Data Storage Alliance, which was founded in 2020 with the aim (again, unlike optical) of promoting compatibility among diverse, proprietary architectures. In June 2022, the Storage Networking Industry Association (SNIA)—a nonprofit technical organization globally recognized and trusted as a storage technology authority—announced that the DNA Data Storage Alliance had joined SNIA as a Technology Affiliate group to accelerate its mission to create an interoperable ecosystem for DNA-based data storage solutions.

The DNA molecule is a physical structure of great beauty and happens to be superbly suited for digital data storage. After all, unique variable-length sequences of only four nucleotides— adenine (A), thymine (T), cytosine (C) and guanine (G)—generate the countless codes that determine the individual chemistry and morphology of all living creatures, from amoeba to whales to human beings. And DNA can endure thousands (perhaps millions) of years in a cool, dark place, with virtually no maintenance or energy costs, and can be accessed by a variety of devices in a future-proof manner with no necessity of data migration. With DNA storage, the data itself is the medium; there are no HDD, SSD, tape or other enterprise-grade media prone to degradation or drive failure and subject to technological change. As long as humans survive, there will be DNA readers.

It makes transcendent sense. Our DNA is, after all, who we are. DNA data storage might comprise an ultimate and enduring record of what we've left behind, in all its garishness and grandeur, during our ephemeral lives.

The eventual space- and power-saving attributes of DNA data storage are still largely unknown, but there is general agreement that they will be formidable, with thousands of times the data density of current technologies packed into "time capsules" half the size of a human finger. In addition to its presumed immense capacity, DNA storage of digital data that is meant to remain permanently unchanged has a crucial advantage compared with current technologies. Once a strand of DNA has been synthesized with unique, encoded binary data, it can be rendered virtually imperishable through thousand-fold replication. The costs of replication will be (theoretically) minimal to nonexistent, pennies or less per copy, and the data can reside anywhere and repose in profuse forms (the tiny "time capsules" being only one proposed "format"), thereby providing, like tape and optical storage, "offline air-gap" security with almost no power consumption. Another consideration is that repatriation/recovery of corporate archive data from the off-premises cloud will be (theoretically) time-consumptive and hugely expensive, whereas even huge amounts of encapsulated on-premises DNA enterprise data can be (theoretically) restored far more quickly and conveniently at far less cost.

DNA data storage, while potentially delivering unprecedented longevity and density with immense resistance to shock, humidity, and temperature extremes, has serious packaging challenges to overcome. It achieves its long life by being hermetically packaged; currently, dehydrated DNA is stored in laser-welded metal capsules to prevent humidity-induced degradation. While this is a sensible and sound solution for storing human DNA samples indefinitely, the use case for enterprise storage will require a solution that makes storage and retrieval more easily accessible while still delivering longevity.

Another aspect of DNA data storage that gives one pause is the fact that processing is water-based instead of electron-based. DNA is manipulated with water as a solvent and with unimaginable numbers of

DNA molecules in solution. Each data set can easily have 100-10,000x redundancy through a simple water-based amplification process. If data recovery is required, a rehydrated sample is extracted from the main solution and destructively read by a DNA sequencer (but only this single copy is destroyed, the other hundreds or thousands of copies remain intact). It appears that enterprise architects will need to rethink how data is stored and retrieved when so many copies are involved.

Inordinate costs and lengthy times to synthesize (hours) and sequence (days) remain daunting challenges. It's a virtual certainty that write times will remain comparatively (and frustratingly) slow, and although new nanopore sequencing technologies have been validated that might substantially reduce retrieval times, there are more doubts than hopes that data access times can be reduced to much less than several hours, but most developers believe the desired goal of $0.001 per gigabyte or less end-user cost for error-free, durable, uncontaminated, deep-but-accessible archive data will be achieved.

The question is, when? Many companies have expressed their willingness to participate in early trials and most companies will consider using DNA data storage in some form. We expect to see a few commercial DNA data storage designs enter an initial phase of production during the next two or three years, but mass-market deployments will not occur until the 2026-2030 time frame.

Costs of Ownership

As with any old-but-enhanced or new storage technology or service, users should demand (whenever possible) case studies from their vendors, and carefully and complexly consider initial acquisition costs and the timing of available capacities in the context of ongoing reliability metrics, performance and power requirements, and the total expense of preserving enterprise data over 5-year, 10-year, 20-year, 50-year, and "indefinite" retention periods.

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 and Twist Bioscience and was written by **John Monroe**

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

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.

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, 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 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