Furthur
Market Research

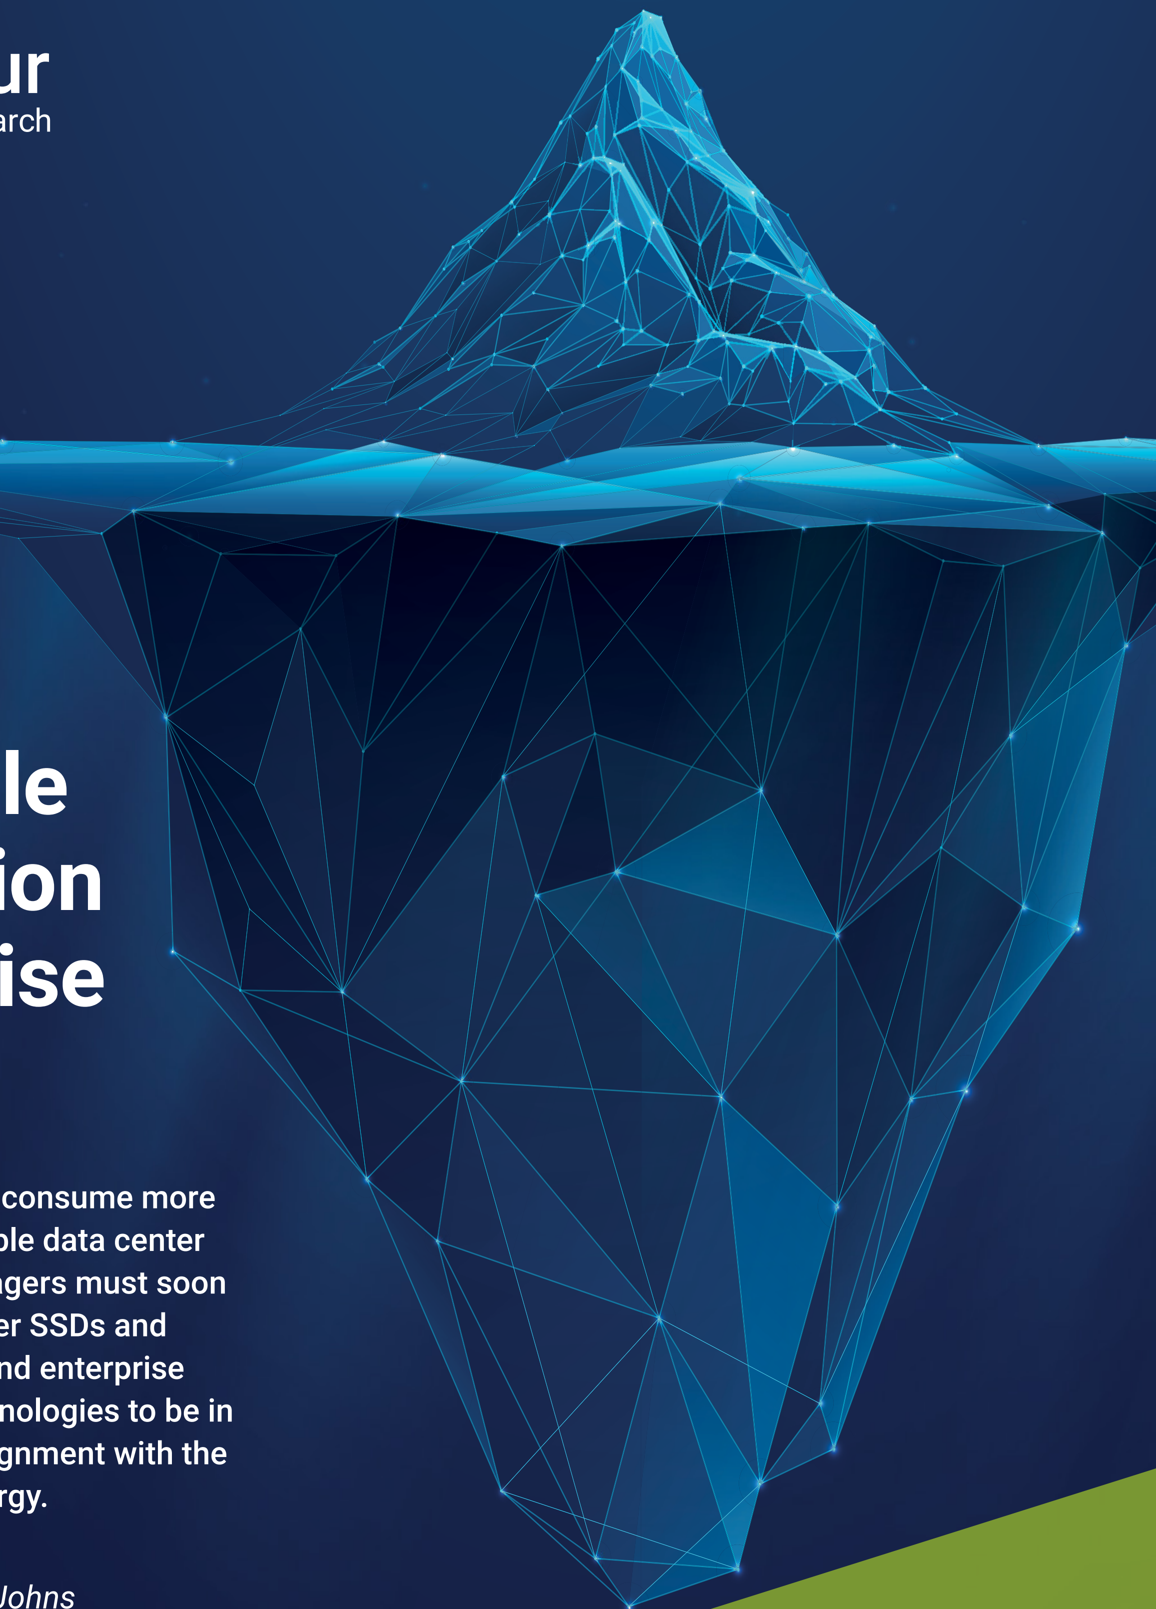

The Sustainable Preservation of Enterprise Data

Enterprise storage will consume more and more of the available data center power budget. IT managers must soon proactively deploy fewer SSDs and HDDs and more tape and enterprise emerging storage technologies to be in crucial and resilient alignment with the total availability of energy.

By John Monroe and Brad Johns

John Monroe | Furthur Market Research
Brad Johns Consulting LLC

January 2024

Table of Contents

Prologue: A Dataverse of Stunning Dimensions

The escalating enormity of our digital dataverse is literally unimaginable but we can with a degree of rigorous certainty determine its current size. At the end of 2022—after a year in which demand declined in unprecedented ways and total enterprise exabyte deliveries expanded by only 2%[1]—the dimensions of the active installed base of enterprise data stored on SSD, HDD, and tape media still grew to 4.8 zettabytes (or 4.8 thousand exabytes, or 4.8 million petabytes), up a staggering 53x over only 91,000 petabytes (or 91 exabytes) in 2010. We estimate the active installed base of enterprise data will exceed 40 zettabytes in 2035, up more than 475x over 2010.

> The billions of people and systems and sensors connected in the global dataverse have generated and will continue to generate immense quantities of data.

We are creating more and more data and deleting less and less of the data we create. Because almost all small and large organizations seem convinced that the only way to remain competitively fit will be to develop agile abilities to derive profit from data, data has become the new "oil" (unlike oil never to be burned but like oil always to be mined for its potential value). "Data is the new oil" has become a cliché, but it should be noted that clichés earn their status as clichés because they are so obviously true. As we enter the AI era, this cliché will become more deeply (and disturbingly) true.

In the context of data center power consumption and our unwillingness to delete any data (no matter how seemingly trivial or insignificant), "data is the new oil" takes on a more sinister dimension.

> Our global digital culture's obsession with and addiction to data could lead to a disproportionate increase in a power consumption metric.

At least 70%, more probably 80%, of all enterprise data is and will continue to be "cool" or "cold" or "frozen," with infrequent access times of minutes to days to weeks to years to decades, with little or no need for the performance of SSDs and HDDs, but with greatly expanding needs for Sustainability, Immutability, and Security (SIS), which SSDs and HDDs can neither cost effectively nor power efficiently fulfill.

In 2022, SSDs accounted for 15.3%, HDDs accounted for 69.5%, and tape accounted for 15.2% of the total enterprise exabytes delivered. The SSD and HDD share of these exabytes has increased while the tape share has declined in recent years, largely due to a costly lack of adequate assessment of actual access-time requirements for workloads in the cool-cold-frozen data layers—an analytical lack which AI technologies may soon alleviate.

> Huge numbers of HDDs and a significant number of SSDs are managing and—without substantial change in purchasing and integration policies—will continue to manage far too many of the cool-cold-frozen workloads at far too great a cost per terabyte while consuming an inordinate share of available energy.

The costs of managing our multi-zettabyte-fold dataverse over increasingly lengthy time periods will continue to swell, and the power demands of enterprise storage will continue to increase as a percentage of the overall data center energy budget. Data center managers must learn to integrate more cost-effective and power-efficient storage technologies.

There are already a multitude of CO2 emission compliance regulations in place throughout the world (with much stricter regulations in Europe) and growing scarcities of total available energy for datacenters in many small communities and metropolitan areas. Healthy ecosystems have become more crucial considerations in all IT purchasing decisions, and many data center managers will soon be forced—by upper-level management edict or by compliance regulations—to use tape and various enterprise emerging technologies as ultra-low-cost, sustainable storage alternatives.

1 *To put this in historical perspective: From 2005-2021, there were no years of <10% enterprise exabyte growth and only two years in the recent past (2017 and 2018, at 19.4% and 19.3%, respectively) of <20% expansion. Even during the 2008-2009 economic meltdown crisis, we saw 42.7% and 35.7% respective annual enterprise exabyte expansions. And even after the "disastrous"/"unprecedented" Q4/2011 effects of the Thai floods, which temporarily reduced enterprise HDD production capabilities by >60%, we saw a 15.9% expansion in enterprise HDD exabytes delivered, and an overall expansion of 27.4%, buoyed by a 43.3% growth in tape exabytes delivered in 2011 (tape accounted for 44.9% of all enterprise exabytes delivered in 2011, its highest % of the total to date).*

Enterprise Data Shipment History and Expansion Estimates, 2010-2035

Enterprise Exabytes

We define enterprise exabytes as the capacities delivered on all enterprise-class SSDs, HDDs, tape, and—in the near future—enterprise emerging storage media. This definition specifically excludes exabyte shipments of consumer-grade SSDs, HDDs, and flash modules delivered to PCs, entertainment devices, cell phones, home video surveillance, and other consumer and industrial applications (such as aircraft and telecom installations), the vast majority of which are already backed up in, and therefore reflected by, the enterprise-grade exabytes serviced by corporate and cloud data centers.

Enterprise Emerging Storage Technologies

Proposed but still nascent emerging technologies include, in alphabetical order: Cerabyte's ceramic nanolayers, Folio Photonics' dynamic multi-layer optical discs, Group47's DOTS (Digital Optical Technology System), and Microsoft's silica. All of these technologies may be available in strategic volume during 2025, but Microsoft's silica will likely be used internally and will not be available externally for commercial consumption. DNA data storage has been generously funded by many companies but will probably have minimal impact prior to 2030. New breeds of tape, as yet uncreated and unspecified, outside of and distinct from the LTO and IBM TS1100 specifications and road maps, will be included in the "Enterprise Emerging Storage" category. In this paper, we refer in Table 1, Table 2, and elsewhere to the total of LTO+IBM+enterprise emerging technology shipments as the "active archive."

Figure 1 depicts the evolving shipment and installed base estimates detailed in Table 1 and Table 2.
Figure 2 depicts the 2023-2035 alternative active installed base scenarios detailed in Table 3.
Table 1 details actual shipments and estimates for the active installed base of enterprise data from 2010 through 2022.
Table 2 details forecast shipments and estimates for the active installed base of enterprise data from 2023 through 2035.
Table 3 details alternative shipment and installed base scenarios from worst-case to unlikely-but-feasible viewpoints.

Figure 1: **Enterprise Data Shipments and the Active Installed Base, 2020-2035**

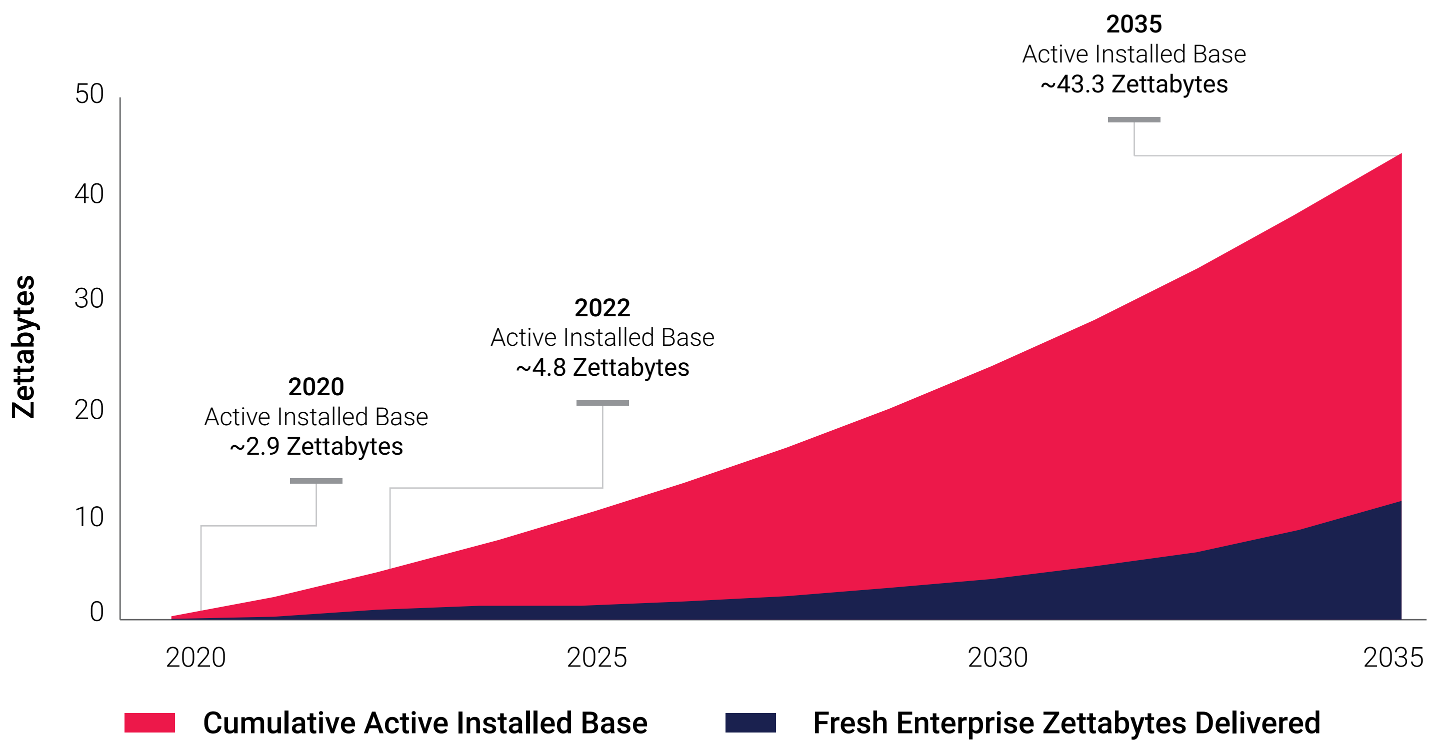

Source: Furthur Market Research (January 2024)

Figure 2: Alternate Active Installed Base Scenarios, 2023-2035

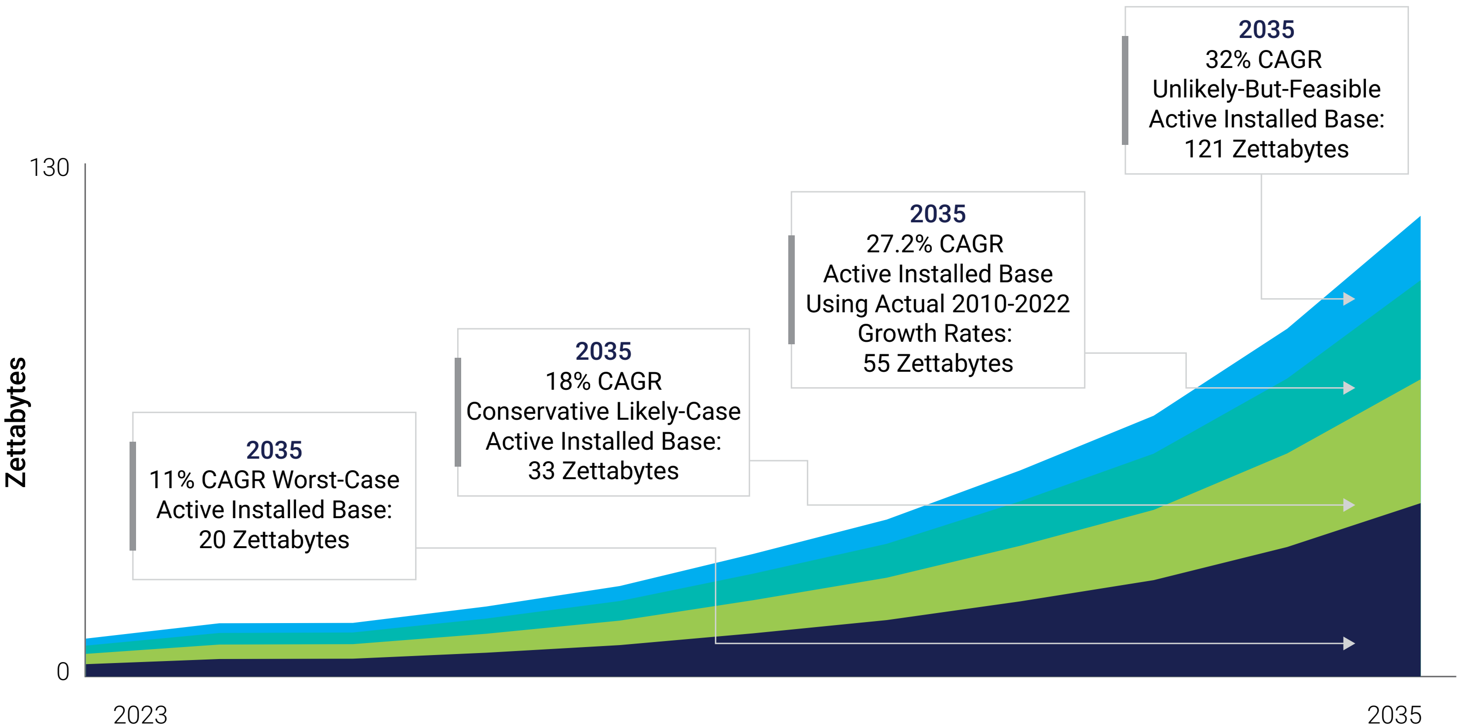

Source: Furthur Market Research (January 2024)

Table 1: Enterprise SSD, HDD, and Tape Actual Shipments and Active Installed Base Estimates in Exabytes, 2010-2022

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	CAGR 2010-2022
Compressed* Enterprise SSD EB	0.2	0.7	1.5	3.8	8.0	13.5	21.4	35.1	61.9	79.7	130.8	179.0	207.6	
YoY Growth %	-	232.8	110.0	153.2	111.3	68.1	58.8	63.8	76.6	28.7	64.1	36.9	16.0	**47.8**
Raw, Uncompressed Enterprise HDD EB	45.2	52.4	65.9	90.6	116.6	157.1	217.9	261.4	333.0	485.7	679.9	959.0	941.7	
YoY Growth %	-	15.9	25.8	37.5	28.7	34.7	38.7	20.0	27.4	45.8	40.0	41.1	-1.8	**28.8**
Total SSD+HDD EB	45.4	53.1	67.4	94.4	124.6	170.6	239.3	296.4	394.9	565.4	810.7	1,138.0	1,149.3	
YoY Growth %	-	16.9	26.9	40.1	32.0	36.9	40.3	23.9	33.2	43.1	43.4	40.4	1.0	**30.9**
Compressed** LTO+IBMTS1100 "Active Archive"Tape EB	30.2	43.3	53.2	61.9	85.6	98.4	123.8	135.9	125.4	134.3	136.1	189.9	206.7	
YoY Growth %	-	43.3	22.8	16.2	38.5	15.0	25.8	9.8	-7.7	7.1	1.4	39.5	8.8	**17.4**
Total Compressed Shipments EB	75.6	96.4	120.6	156.3	210.2	269.0	363.1	432.4	520.3	699.7	946.8	1,327.9	1,356.0	
YoY Growth %	-	27.4	25.0	29.6	34.6	34.0	29.7	19.4	19.3	34.5	33.6	40.3	2.1	**27.2**
SSD EB % ofTotal Annual Shipments	0.3	0.7	1.2	2.4	3.8	4.8	5.9	8.0	11.9	11.4	13.8	13.5	15.3	
HDD EB % ofTotal Annual Shipments	59.8	54.4	54.7	58.0	55.5	55.8	59.7	59.9	64.0	69.4	71.8	72.2	69.5	
SSD+HDD EB % ofTotal Annual Shipments	60.1	55.1	56.0	60.5	59.3	60.6	65.5	68.0	75.9	80.8	85.6	85.7	84.8	
Tape EB % ofTotal Annual Shipments	39.9	44.9	44.1	39.6	40.7	34.9	33.9	31.2	24.1	19.2	14.4	14.3	15.2	
Active Installed Base EB	91.0	180.5	291.1	430.2	613.9	819.9	1,088.8	1,404.0	1,768.6	2,258.1	2,923.2	3,885.8	4,805.8	
YoY Growth %	-	98.3	61.3	47.8	42.7	33.6	32.8	29.0	25.9	27.7	29.5	32.9	23.7	**39.2**
Active Installed Base xChange Relative to 2010	-	2.0	3.2	4.7	6.7	9.0	12.0	15.4	19.4	24.8	32.1	42.7	52.8	

Source: Furthur Market Research (January 2024)

Table 2: Enterprise SSD, HDD, Tape, and Emerging Storage Forecast Shipments and Active Installed Base Estimates in Exabytes, 2023-2035

	2023	2024	2025	2026	2027	2028	2029	2030	CAGR 2023 or 2025 to 2030	2031	2032	2033	2034	2035	CAGR 2031-2035
Compressed* Enterprise SSD EB	129.9	149.5	198.2	258.5	327.3	436.9	561.0	735.6		930.5	1,165.9	1,442.2	1,753.7	2,074.7	
YoY Growth %	-37.4	15.1	32.6	30.4	26.6	33.5	28.4	31.1	**28.1**	26.5	25.3	23.7	21.6	18.3	**22.2**
Raw, Uncompressed Enterprise HDD EB	1,038.9	1,357.9	1,700.0	2,203.2	2,833.4	3,436.9	4,275.5	5,066.4		4,645.9	3,911.9	3,203.8	2,579.1	1,833.7	
YoY Growth %	10.3	30.7	25.2	29.6	28.6	21.3	24.4	18.5	**25.4**	-8.3	-15.8	-18.1	-19.5	-28.9	**-20.7**
Total SSD+HDD EB	1,168.8	1,507.3	1,898.3	2,461.7	3,160.7	3,873.8	4,836.5	5,802.0		5,576.4	5,077.8	4,646.1	4,332.8	3,908.4	
YoY Growth %	1.7	29.0	25.9	29.7	28.4	22.6	24.9	20.0	**25.7**	-3.9	-8.9	-8.5	-6.7	-9.8	**-8.5**
Compressed** LTO+IBMTS1100 "Active Archive" Tape EB	228.4	264.7	314.2	384.2	459.5	591.8	757.0	948.5		1,167.6	1,464.2	1,751.1	2,054.1	2,460.8	
YoY Growth %	10.5	15.9	18.7	22.3	19.6	28.8	27.9	25.3	**22.6**	23.1	25.4	19.6	17.3	19.8	**20.5**
Enterprise Emerging Storage (2025 Onward)			18.8	83.2	123.6	263.5	488.9	735.0		998.2	1,384.5	1,874.6	2,753.7	3,904.8	
YoY Growth %			-	343.4	48.5	113.2	85.6	50.3	**108.3**	35.8	38.7	35.4	46.9	41.8	**40.6**
Total "Active Archive" Storage Opportunity	228.4	264.7	332.9	467.4	583.1	855.3	1,245.9	1,683.5		2,165.8	2,848.6	3,625.7	4,807.8	6,365.6	
YoY Growth %	10.5	15.9	25.8	40.4	24.7	46.7	45.7	35.1	**38.3**	28.6	31.5	27.3	32.6	32.4	**30.9**
Total Compressed Shipments EB	1,397.1	1,772.0	2,231.2	2,929.1	3,743.7	4,729.1	6,082.3	7,485.5		7,742.2	7,926.4	8,271.7	9,140.6	10,274.0	
YoY Growth %	3.0	26.8	25.9	31.3	27.8	26.3	28.6	23.1	**27.1**	3.4	2.4	4.4	10.5	12.4	**18.1**
SSD EB % of Total Annual Shipments	9.3	8.4	8.9	8.8	8.7	9.2	9.2	9.8		12.0	14.7	17.4	19.2	20.2	
HDD EB % of Total Annual Shipments	74.4	76.6	76.2	75.2	75.7	72.7	70.3	67.7		60.0	49.4	38.7	28.2	17.8	
SSD+HDD EB % of Total Annual Shipments	83.7	85.1	85.1	84.0	84.4	81.9	79.5	77.5		72.0	64.1	56.2	47.4	38.0	
Tape EB % of Total Annual Shipments	16.3	14.9	14.1	13.1	12.3	12.5	12.4	12.7		15.1	18.5	21.2	22.5	24.0	
Enterprise Emerging EB % of Total Annual Shipments			0.8	2.8	3.3	5.6	8.0	9.8		12.9	17.5	22.7	30.1	38.0	
"Active-Archive" Storage (Tape+Emerging) % of Total EB			14.9	16.0	15.6	18.1	20.5	22.5		28.0	35.9	43.8	52.6	62.0	
Active Installed Base EB	5,682.6	6,754.9	8,039.3	9,640.5	12,028.3	15,360.2	19,670.5	24,924.9		29,737.9	33,920.6	37,463.3	40,521.6	43,310.1	
YoY Growth %	18.2	18.9	19.0	19.9	24.8	27.7	28.1	26.7	**23.5**	19.3	14.1	10.4	8.2	6.9	**18.4**
Active Installed Base xChange Relative to 2010	62.4	74.2	88.3	105.9	132.2	168.8	216.2	273.9		326.8	372.8	411.7	445.3	475.9	

Source: Furthur Market Research (January 2024)

Table 3: Alternate Shipment and Installed Base Scenarios in Exabytes, 2023-2035

	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035
Alternate 2023-2035 Shipment Scenarios: Worst-Case Viewpoints													
Total Shipped Enterprise EB Expanding at 11%/Year 2023-2035	1,397.1	1,550.8	1,721.4	1,910.7	2,120.9	2,354.2	2,613.2	2,900.7	3,219.7	3,573.9	3,967.0	4,403.4	4,887.8
Active Installed Base EB at 11% Annual Shipment Expansion	5,682.6	6,533.7	7,308.3	7,891.2	8,656.1	9,613.2	10,675.6	11,854.9	13,163.9	14,616.8	16,229.6	18,019.8	20,006.9
Alternate 2023-2035 Shipment Scenarios: Conservative Likely-Case Viewpoints													
Total Shipped Enterprise EB Expanding at 18%/Year 2023-2035	1,397.1	1,648.6	1,945.4	2,295.5	2,708.7	3,196.3	3,771.6	4,450.5	5,251.6	6,196.9	7,312.3	8,628.5	10,181.7
Active Installed Base EB at 18% Annual Shipment Expansion	5,682.6	6,631.5	7,630.1	8,597.7	9,950.4	11,749.6	13,872.6	16,377.7	19,333.8	21,876.7	24,877.3	28,418.0	32,596.0
Alternate 2023-2035 Shipment Scenarios: Using Actual 2010-2022 Growth Rates													
Total Shipped Enterprise EB Expanding at 27.2%/Year 2023-2035	1,397.1	1,777.1	2,254.0	2,838.1	3,725.8	4,762.0	6,015.4	7,736.7	9,521.6	9,848.1	10,082.4	10,521.7	11,626.9
Active Installed Base EB at 27.2% Annual Shipment Expansion	5,682.6	6,760.1	8,067.3	9,577.4	11,947.3	15,312.2	19,550.4	25,033.1	31,716.7	39,310.7	45,667.3	51,426.9	55,317.1
Alternate 2023-2035 Active Installed Base Scenarios: Unlikely-But-Feasible Viewpoints													
Total Shipped Enterprise EB Expanding at 32%/Year 2023-2035	1,397.1	1,844.2	2,434.3	3,213.3	4,241.6	5,598.9	7,390.6	9,755.6	12,877.3	16,998.1	22,437.5	29,617.5	39,095.1
Active Installed Base EB at 32% Annual Shipment Expansion	5,682.6	6,827.1	8,314.7	10,200.1	13,085.7	17,287.5	22,833.9	30,155.1	39,819.1	52,575.6	69,414.2	91,641.1	120,980.6

Note: 4Q23 forecasts are used as common starting points for all scenarios.

Source: Furthur Market Research (January 2024)

Notes relevant to Tables 1, 2, and 3:

- **SSD capacities reflect an approximate 5x compression ratio, but only for approximately 5% of all enterprise SSD EBs 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 software is the norm. Enterprise SSD CAGR in Table 1 is calculated for 2015-2022 to provide meaningful comparisons with HDD and Tape, since 2010-2014 enterprise SSD shipments were so minimal.*
- *HDD capacities are raw/uncompressed, since so few enterprise HDDs utilize any form of data compression.*
- ***Tape capacities include both LTO and IBM TS1100 shipments and reflect a global average of 2.5x data compression.*
- *Enterprise optical shipments have remained 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 or the current active installed base. That said, there should be huge opportunities for what we are now referring to collectively as "Enterprise Emerging Storage" technologies to play major roles in future enterprise markets, as indicated in our 2025-2035 growth estimates. The emerging "Total 'Active Archive' Storage Opportunity" delineated in Table 2 is the sum of all LTO+IBM+enterprise emerging technology shipments.*
- *SSD, HDD, and LTO+IBM CAGR is calculated in Table 2 from 2023-2030 and 2031-2035; enterprise emerging technology CAGR is calculated in Table 2 from 2025-2030 and 2031-2035.*
- *We estimate the active installed base of total enterprise capacities was 91 exabytes in 2010 and will grow to exceed 40 zettabytes in 2035. 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 installed base of the prior year, and we repeat this cycle through 2035.*

Astonishing Growth in the Active Installed Base

The active installed base exceeded one zettabyte in 2016 and grew to 4.8 zettabytes in 2022. We are squarely in the midst of the zettabyte era, with the active installed base likely exceeding 40 zettabytes and perhaps growing to 50 zettabytes or more in 2035. Even in a worst-case 11%-per-year growth scenario, the active installed base in 2035 will have expanded to more than 20 zettabytes, more than 200x over 2010. In a 32% growth scenario—however unlikely, given recent results, but it should still be considered and not immediately dismissed—the active installed base might expand to more than 100 zettabytes in 2035.

Will the Past Be Prologue or Will History Be Bunk?

Compared with our prior forecasts, we have significantly decreased our estimates of SSD and HDD exabytes shipped and accelerated the potential impact of tape and emerging technologies. A key question remains: Will the "unprecedented" declines of 2022-2023 be an anomaly or a harbinger of a "new normal" landscape of evolving demand (reflected in our worst-case scenario)? Growth rates will surely fluctuate, but we do not think the 2.1% actual and ~3% forecast growth rates of 2022 and 2023—with hyperscale and other large enterprise customers delaying or cancelling orders to an inordinate degree and digesting inventory 3Q22-4Q23—will be repeated with any degree of cyclic frequency 2023-2035. For example, in December 2023, we saw dramatic surges in NAND demand and pricing, and widespread product allocations in multiple markets. Total enterprise SSD ASP/unit is now expected to increase from $184 in 2023 to $312 in 2024.

Assumptions

Exabyte Shipments

- After the unprecedented -80% price erosions and -37% downturns in demand 2Q22-4Q23, enterprise SSD exabytes will grow consistently through 2035, but at ameliorated rates, because a) technology transitions will prove to be slower and more costly 2023-2035, b) ASPs/TB will remain relatively uncompetitive—at least a 2-3x SSD:HDD and a 10x SSD:Tape/Emerging differential through 2035, and c) limited fab capex investments will limit production and allocation capabilities to less than 1 enterprise zettabyte in 2030. Given current ASP and allocation trends, combined with demand for NAND in PCs and mobile phones, and escalating costs to produce sufficient NAND exabytes for all markets, it is hard to imagine that, irrespective of demand, enterprise SSD capacity shipments can far exceed two zettabytes in 2035, even if the NAND makers allocate more than 50% of their production solely to enterprise SSDs. Our current 2023 estimates for raw NAND allocation are 108.2 exabytes going to enterprise SSDs—which, compressed, will service 129.9 exabytes of enterprise data demand—or 13.5% of the total 803.7 raw exabytes that will be produced for all markets. The percentage of raw NAND allocated to enterprise SSDs exceeded 20% in 2018, 2020, 2021, and 2022, but will remain below 20% through 2027.
- HDD shipments will peak at ~5.1 zettabytes in 2030 and decline 2031-2035. Due to slowing areal density growth, modulating ASP declines, disciplined market management in the face of competing technologies, and restrained infrastructure investments, it is unlikely that the HDD makers will ever be able to deliver much more than ~5 zettabytes/year.
- Despite 2031-2035 declines, HDD exabyte deliveries in 2035 will greatly exceed the HDD exabytes delivered in 2022 (1.83 vs. 0.94 zettabytes).
- Enterprise SSD exabyte deliveries will exceed enterprise HDD exabyte deliveries beginning in 2035.
- Tape and enterprise emerging technology shipments will display consistent growth through at least 2035 and will exceed combined SSD+HDD exabyte deliveries in 2034.

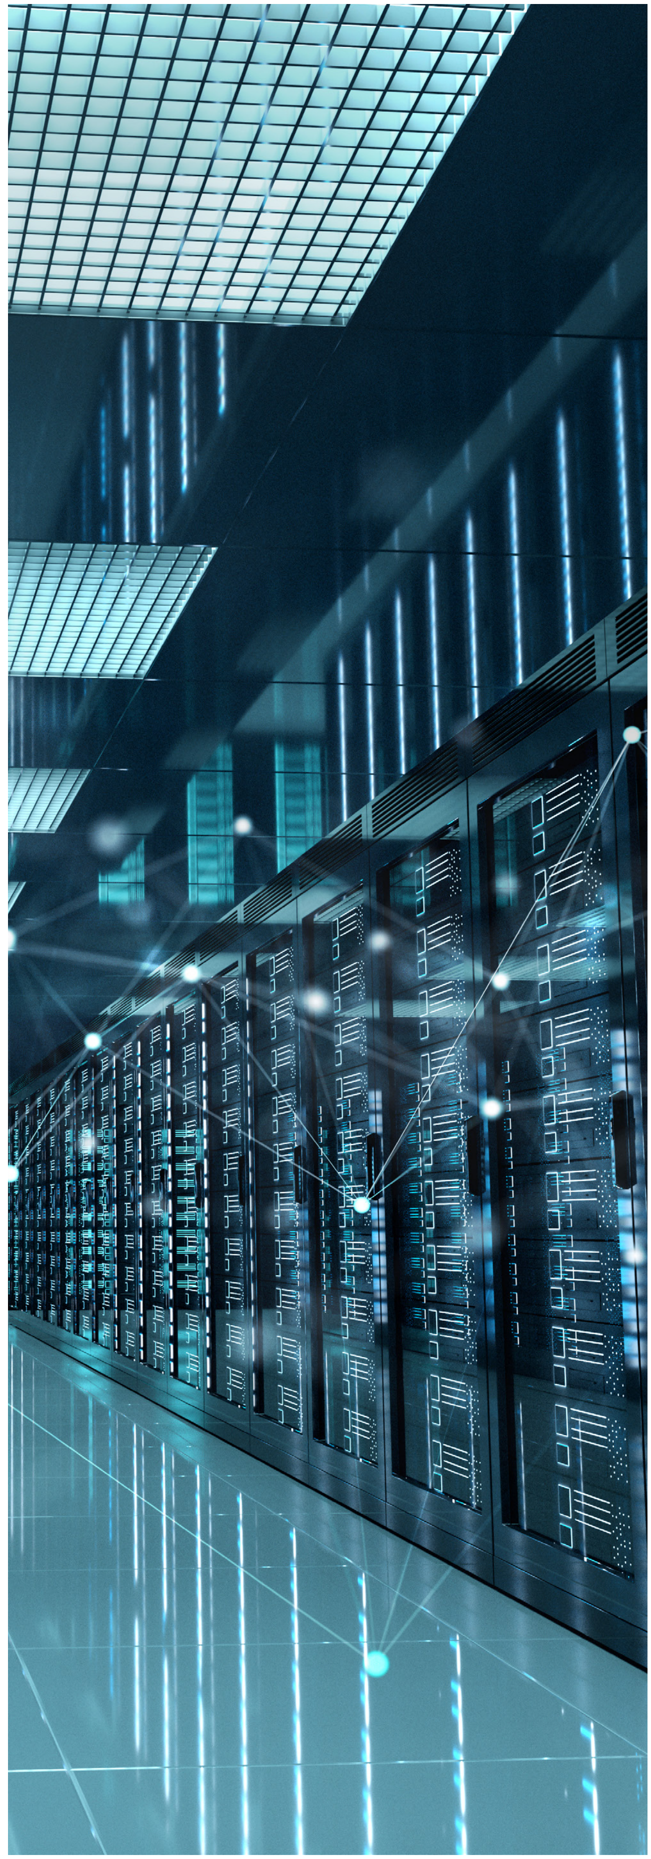

ASP Considerations and Long-Term Sustainability

- At least 70% of all enterprise data will be "cool" or "cold" or "frozen" (see Figure 2 below), with access frequencies spanning minutes to days to weeks to years to decades, with little or no need for the performance of SSDs and HDDs, but with greatly expanding needs for Sustainability, Immutability, and Security (SIS), which SSDs and HDDs can neither cost effectively nor power efficiently fulfill.
- In the near term, tape and enterprise emerging technology infrastructures will consume ~99% less power than SSD and HDD infrastructures.
- Tape as a percentage of enterprise exabyte deliveries has shrunk from 35% in 2015 to 15% in 2022, largely due to a costly lack of adequate focus on the actual access-time requirements of workloads in the cool-cold-frozen data layers; enterprise SSD+HDD exabytes comprised 85% of shipments in 2022, up from 65% in 2015.
- The lowest enterprise ASPs/TB for products purchased by the largest direct OEMs and end user customers of the drive makers during 2023 were $72.10 for SSDs, $9.90 for HDDs, and $5.20 for tape. SSD and HDD ASPs are per-drive estimates. Tape ASP is determined by lowest costs for robotics, drives, and media; a maximum of 128 drives per library; and an average of 120 cartridges per drive.
- By 2035, we believe it will be difficult for SSD and HDD ASPs to profitably fall to less than $10.00/TB and $6.00/TB, respectively, while tape and enterprise emerging technology prices can profitably fall to less than $1.00/TB.
- Because of cost and energy requirements, CTOs, CIOs, CSOs (Chief Sustainability Officers), and CFOs will increasingly recognize the growing necessity of integrating tape and enterprise emerging technologies in their storage infrastructures.
- Enterprise emerging storage will display the greatest growth 2025-2030 and 2031-2035. Because of increasing appreciation of their long-term strategic storage virtues, enterprise emerging-technology exabytes delivered will exceed tape exabytes delivered in 2033.
- Combined annual long-term, "active archive" storage opportunities will exceed 1.6 zettabytes in 2030 and will grow to more than 6.4 zettabytes in 2035.
- These markets will not be defined by binary, elegant "either/or" scenarios of storage infrastructure choices (for example, either SSD or tape, hot or not, as some analysts have predicted), but rather by complicated "both/and/and/and" scenarios of diverse enterprise technologies used in concert and conjunction with each other.

Access Frequency of Enterprise Data

In our current delineations of enterprise data in terms of access frequency, we have divided the evolving enterprise data layers into the following segments: hot (nanoseconds to milliseconds), warm (milliseconds to seconds), cool (minutes to 24 hours), cold (days to weeks), and frozen (weeks to years). These access frequencies have nothing to do with “time to first byte” specifications of various technologies, but rather the time interval between when the data is initially written and when a user requests access to that data. Most enterprise data becomes cold or frozen after 60 days, and is often kept indefinitely.

The amount of data being delivered to all the layers will continue to expand, but the data distribution percentages will change. Our 2020-2035 most likely assumption is that delivery of exabytes destined to be integrated in the hot and warm data layers will shrink from ~25% (8% hot, 17% warm) to ~20% (7% hot, 13% warm) of the total installed base. Meanwhile, the cool and cold data layers remain fairly constant at ~20% and ~25% of the total installed base, respectively. The frozen data layer will expand from ~30% to ~35% of the annual total. Figure 3 depicts this assumption. Figure 4 updates our storage pyramid to reflect these new delineations of enterprise data.

Figure 3: **Delineations of Enterprise Data Based on Access Frequency**

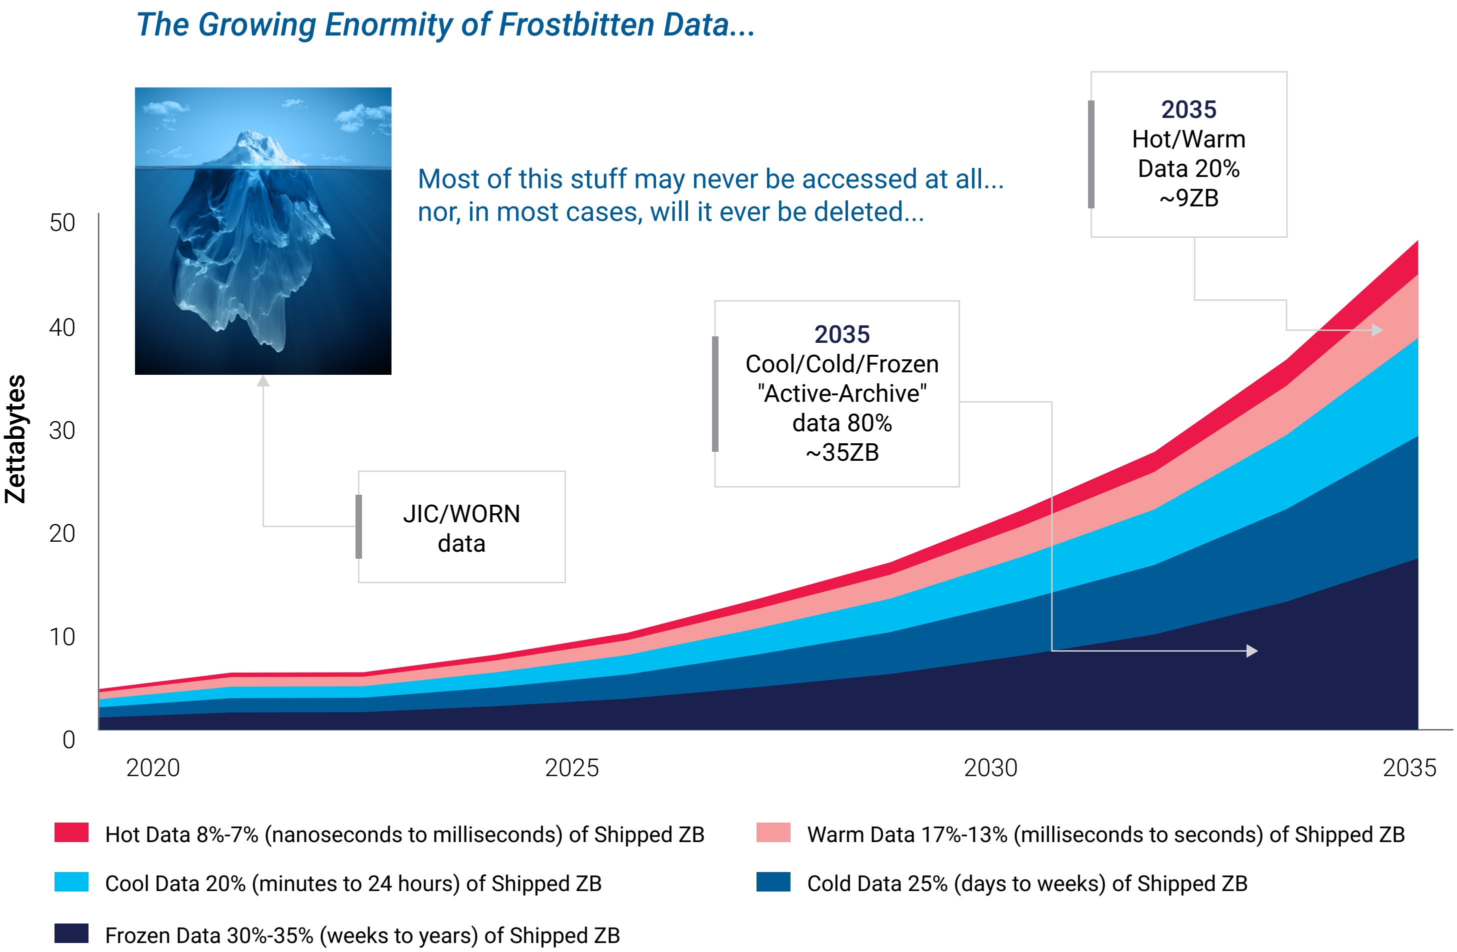

Source: Furthur Market Research (January 2024)

Figure 4: The Evolving Storage Pyramid

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

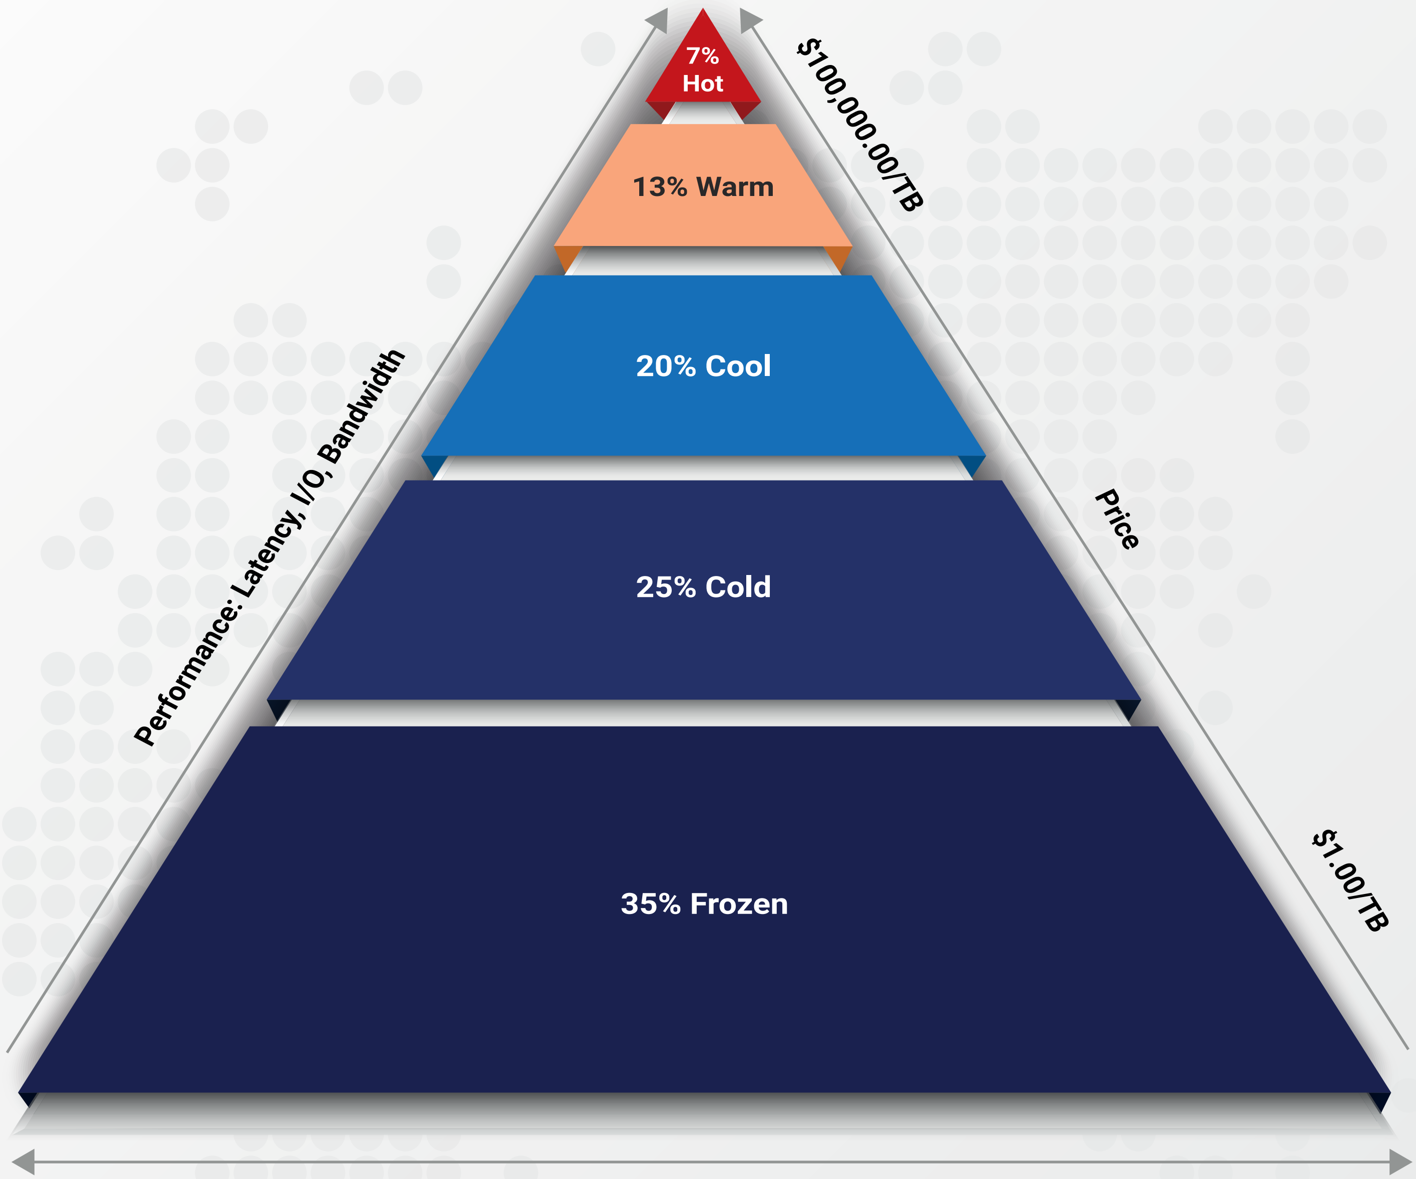

A majority of the cold/frozen layers may be JIC (Just in Case) or WORN (Write Once Read Never) data, which may never be accessed at all–nor, in most cases, will it ever be deleted.

Notes:

- *The hot data layer will be comprised of SRAM/DRAM, storage class memory (SCM), and SSD technologies; SSDs will continue to service more than 95% of this layer.*
- *HDDs currently service more than 10% of the warm data layer but SSDs will service almost 100% of this layer after 2030.*
- *HDDs currently service ~95% of the cool and cold data layers, but active archive technologies will be increasingly integrated in these layers.*
- *Tape technologies currently service only ~40% of the frozen data layer; HDDs currently service the majority of this layer.*
- *There will be some SSDs used as high-speed caches in the cool-cold-frozen data layers, but the majority of these layers will continue to be serviced by HDDs through 2030; after 2030, HDD shipments will shrink while tape and enterprise emerging technologies expand to handle more than 50% of all enterprise data.*
- *See relative distributions of these technologies in Table 1 and Table 2.*

Source: *Furthur Market Research (January 2024)*

Will All Data Simply Be Hot or Not?

In an alternate scenario, with the advent of more strictly enforced corporate archive and access rules, and the growing need to conduct AI/ML business at the speed of flash, it is also possible that, in many data centers, ~30% of the data will be classified as hot in 2035, while the warm and cool and cold data layers diminish to insignificance, and the frozen data layer grows to ~70% of the total—there will be no fine distinctions, either the data is hot, or it's not.

But we think this scenario is unlikely because, as we said above, these markets will not be defined by binary, elegant "either/or" scenarios of storage choices (for example, either SSD or tape, hot or not, as some have predicted), but rather by complicated "both/and/and/and" scenarios of diverse technologies used in concert and conjunction with each other.

In any case, we estimate that annual shipments of enterprise storage capacity destined to manage only cold or frozen data—with access frequencies spanning days to years—will approach 900 exabytes in 2023 and expand in 2035 to ~7.3 zettabytes, while the active installed base of cold and frozen data grows from ~3.1 zettabytes in 2023 to ~25.9 zettabytes in 2035. The active installed base of all cool-cold-frozen data will grow to ~34.6 zettabytes in 2035.

A way to frame these delineations in several real-life examples that support our assumptions regarding the cool-cold-frozen layers, can be seen in the results of our recent interviews with seven senior IT managers responsible for the administration of large—50 petabyte-to-500 petabyte—databases.

- All of these data center managers specified "indefinite" retention periods for the vast majority of their data, but they fear that the rising sustainability costs of preserving their data for many years or "indefinitely" will become prohibitive.
- For all these managers, data immutability was a crucial issue—all aspects of the original data absolutely must remain unchanged. Many of these IT managers classified their data as 100% "cold," but it could (unpredictably) become "hot" at any time depending on data access requests—in other words, 100% of their data is an "active archive."
- Several enterprise IT managers with whom we spoke stated that an exacerbated problem with any data deletion is establishing generally agreed-upon ground rules. (Many kinds of data, such as X-ray or MRI patient data, cannot be deleted at will, but other kinds of data, such as aging emails, can be deleted by corporate rule.) When these managers 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, because 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—"data is the new oil"—for some unspecified, but critical, future purpose.

Shifts in Data Center Power Consumption

Figure 5: Estimated Percentages of Data Center Power Use, 2020 and 2035

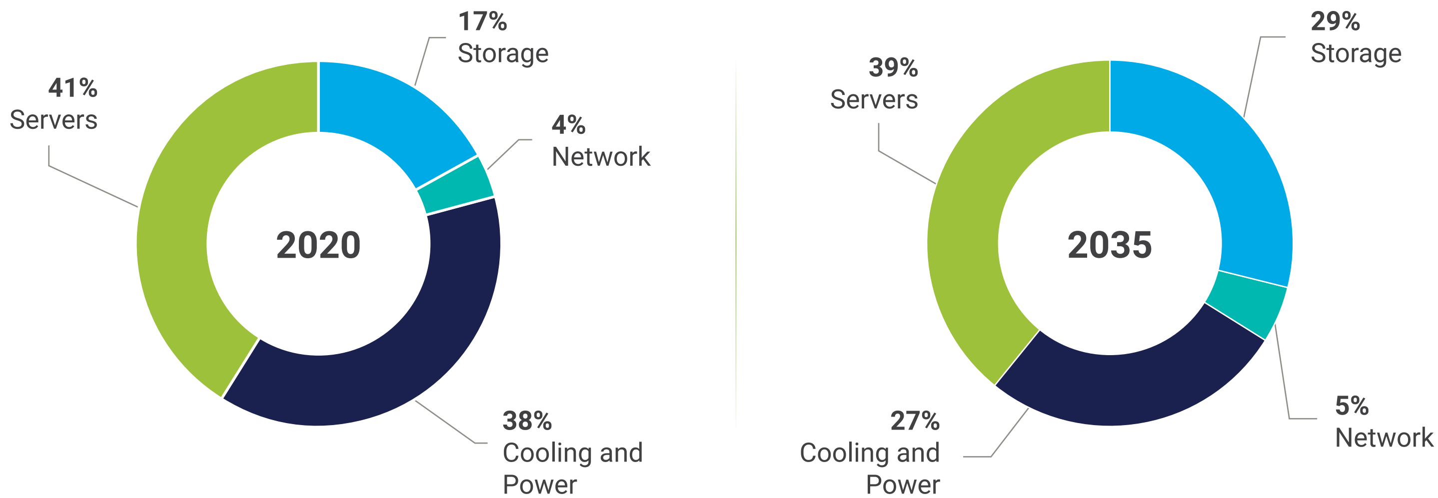

Source: Furthur Market Research (January 2024)

To date, servers and cooling systems have far exceeded storage power requirements in enterprise data centers. However, because of the rapid expansion of digital data needs, in recent years storage as a percent of enterprise data center power needs has expanded and could grow to account for ~29% of the power needs of enterprise data centers in 2035, up from 17% in 2020.

While per-terabyte power consumption will decline in concert with advances in areal density technologies, per-drive power draw will not; in fact, per-drive power draw will increase in concert with increasing shipments of enterprise-grade, NVMe-interface SSDs (see Tables 4 and 5 in Section IV). This will result in a growing imbalance between increases in storage power needs and decreases of available energy.

Figure 6 delineates our estimate that—without substantial change in purchasing and integration policies—storage power consumption will increase by 50x, 2010-2050.

Figure 6: 2050 Enterprise Storage Power Consumption Relative to 2010

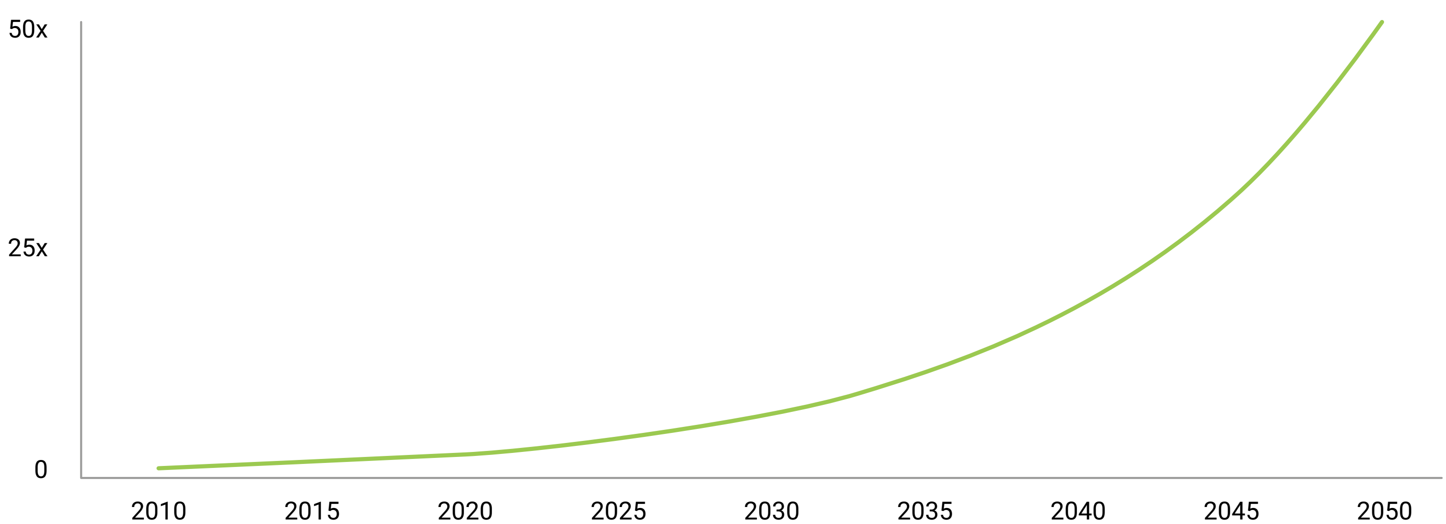

It is clear that enterprise storage requirements are increasing faster than power efficiency gains, which means that enterprise storage infrastructures will consume greater and greater amounts of the data center power budget, creating inflection points where diverse breakthrough technologies will be required to change the shape of this curve.

IV Estimated Raw Power Draw for the Essential Building Blocks of Enterprise Storage

Before we get into our complex and intricate TCO calculations in Section V, we'd like to simply state a few arresting facts regarding the raw megawatts that—regardless of any rack and cooling and other storage system power requirements, and irrespective of per-drive capacities and total zettabytes in the active installed base—will be needed just to drive the bare bones enterprise SSDs and HDDs.

Because power draw for configured server direct-attached (DAS) or external controller-based (ECB) storage systems can vary so widely by vendor and product type, we will focus on only one key metric here: power draw watts per SSD or HDD unit. Our aim is to clarify and simplify without being simplistic.

Figure 7 depicts relative amounts of HDD, SSD, and active archive power draw from 2020 through 2035.
Table 4 details an analysis of power draw for bare bones SSDs and HDDs from 2020 through 2025.
Table 5 extends this analysis from 2026 through 2035.
Comparative estimates for the installed base of tape and other "active archive" technologies are included in Table 4 and Table 5.

We have performed a scrupulous analysis of tape system power draw requirements and used this analysis to estimate total power draw for all tape and emerging technology systems based on the total "active archive" exabyte shipments and forecasts detailed in Table 1 and Table 2. This estimate likely overstates the total power draw requirements for these systems, since all enterprise emerging technologies must have manifestly lower power draw requirements when compared with tape in order to gain any market acceptance at scale.

Figure 7: Annual SSD, HDD, and Total "Active Archive" Power Draw, 2020-2035

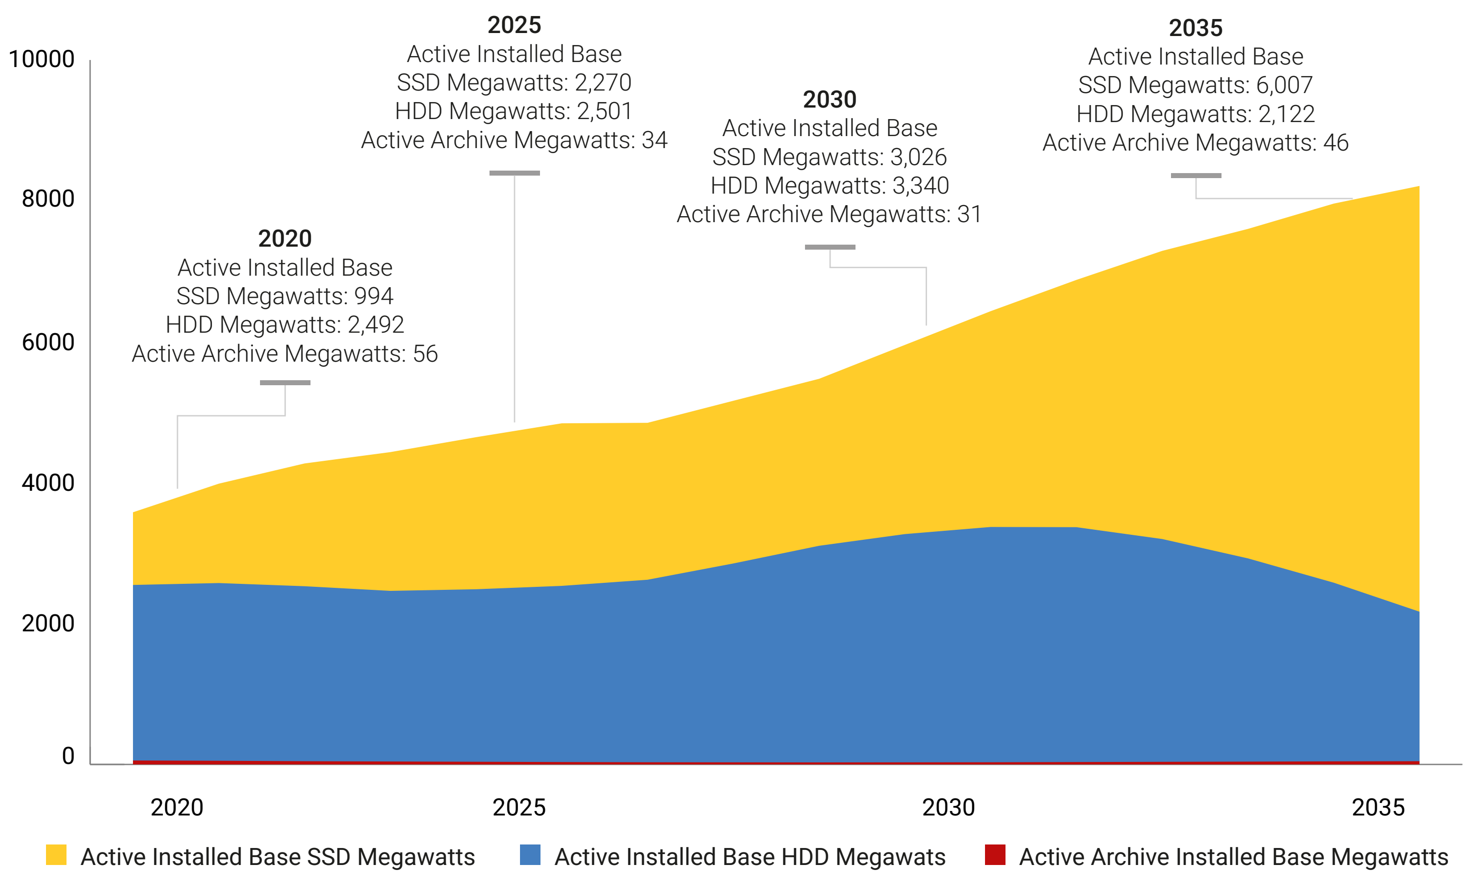

Note the steep decline in HDD power draw as shipments decline after 2030. The active archive megawatts are barely discernible in this format.

Source: Furthur Market Research and Brad Johns Consulting (January 2024)

Table 4: **Enterprise SSD and HDD Estimated Per-Year Unit-Volume Megawatt Power Consumption, 2020-2025**

	Watts	2020	2021	2022	2023	2024	2025	2020-2025 Totals
Total Enterprise SSD K Units		54,240	65,609	66,204	42,365	40,570	43,890	**312,878**
Active Installed Base K Units		162,781	209,657	250,115	262,794	268,988	258,638	**NA**
Total NVMe Enterprise SSD K Units		32,584	48,178	47,265	29,849	31,706	36,730	**252,611**
NVMe % of Total		60	73	71	70	78	84	**73**
Typical Operating Power Watts Per NVMe Drive	11							
Typical Megawatt Power Draw for New Units		358	530	520	328	349	404	**2,779**
Active Installed NVMe Base Megawatt Power Draw		806	1,153	1,458	1,706	1,923	2,134	**9,633**
Total SATA/SAS Enterprise SSD K Units		21,656	17,431	18,939	12,516	8,864	7,160	**93,697**
SATA/SAS % of Total		40	27	29	30	22	16	**27**
Typical Operating Power Watts Per SATA/ SAS Drive	2.1							
Typical Megawatt Power Draw for New Units		45	37	40	26	19	15	**197**
Active Installed SATA/SAS Base Megawatt Power Draw		188	220	247	226	198	136	**1,252**
Total Enterprise HDD K Units		71,851	77,536	67,591	68,390	79,448	86,027	**450,843**
Active Installed Base K Units		347,371	358,608	361,354	355,914	364,816	378,992	**NA**
Total SATA Enterprise HDD K Units		50,297	60,538	55,724	60,049	72,695	80,873	**380,175**
SATA % of Total		70	78	82	88	91	94	**84**
Typical Operating Power Watts Per SATA Drive	6.3							
Typical Megawatt Power Draw for New Units		317	381	351	378	458	509	**2,395**
Active Installed SATA Base Megawatt Power Draw		1,357	1,536	1,707	1,749	1,885	2,077	**10,311**
Total SAS Enterprise HDD K Units		21,554	16,998	11,867	8,341	6,753	5,154	**70,668**
SAS % of Total		30	22	18	12	9	6	**16**
Typical Operating Power Watts Per SAS Drive	8.6							
Typical Megawatt Power Draw for New Units		185	146	102	72	58	44	**608**
Active Installed SAS Base Megawatt Power Draw		1,135	987	777	673	565	424	**4,561**
Annual SSD Megawatt Power Draw		404	567	560	355	367	419	**2,671**
Annual HDD Megawatt Power Draw		502	528	453	450	516	554	**3,003**
Installed Base SSD Megawatt Power Draw		994	1,373	1,705	1,932	2,121	2,270	**10,394**
Installed Base HDD Megawatt Power Draw		2,492	2,523	2,484	2,422	2,449	2,501	**14,872**
Installed Base SSD + HDD Megawatt Power Draw		3,486	3,896	4,189	4,354	4,570	4,771	**25,267**
Installed Base Total "Active Archive" Megawatt Power Draw		46	40	36	33	30	26	**211**
HDD + SSD vs "Active Archive" Power Draw Ratio		76	97	116	132	152	183	**120**

Source: Furthur Market Research and Brad Johns Consulting (January 2024)

Table 5: Enterprise SSD and HDD Estimated Per-Year Unit-Volume Megawatt Power Power Draw, 2026-2035

	Watts	2026	2027	2028	2029	2030	2026-2030 Totals	2031	2032	2033	2034	2035	2031-2035 Totals
Total Enterprise SSD K Units		46,856	47,820	55,789	63,035	73,181	**286,681**	83,744	95,667	109,783	123,454	133,440	**546,088**
Active Installed Base K Units		239,885	221,501	234,925	257,390	286,681	**NA**	323,569	371,416	425,410	485,829	546,088	**NA**
Total NVMe Enterprise SSD K Units		42,073	44,074	52,708	61,166	72,336	**272,358**	83,744	95,667	109,783	123,454	133,440	**546,088**
NVMe % ofTotal		90	92	94	97	99	**95**	100	100	100	100	100	**100**
Typical Operating Power Watts Per NVMe Drive	11												
Typical Megawatt Power Draw for New Units		463	485	580	673	796	**2,996**	921	1,052	1,208	1,358	1,468	**6,007**
Active Installed NVMe Base Megawatt Power Draw		2,085	2,193	2,274	2,605	2,996	**12,152**	3,452	4,045	4,633	5,344	6,007	**23,481**
Total SATA/SAS Enterprise SSD K Units		4,783	3,746	3,081	1,869	845	**14,323**	0	0	0	0	0	**0**
SATA/SAS % ofTotal		10	8	6	3	1	**5**	0	0	0	0	0	**0**
Typical Operating Power Watts Per SATA/SAS Drive	2.1												
Typical Megawatt Power Draw for New Units		10	8	6	4	2	**30**	0	0	0	0	0	**0**
Active Installed SATA/SAS Base Megawatt Power Draw		106	79	59	43	30	**317**	20	8	7	0	0	**35**
Total Enterprise HDD K Units		96,974	107,818	110,975	108,939	103,760	**528,466**	96,923	81,609	66,838	53,604	38,113	**337,087**
Active Installed Base K Units		398,222	438,449	481,034	510,525	528,258	**NA**	528,207	501,998	457,861	402,527	336,879	**NA**
Total SATA Enterprise HDD K Units		93,863	105,928	110,210	108,939	103,760	**522,700**	96,923	81,609	66,838	53,604	38,113	**337,087**
SATA % ofTotal		97	98	99	100	100	**99**	100	100	100	100	100	**100**
Typical Operating Power Watts Per SATA Drive	6.3												
Typical Megawatt Power Draw for New Units		591	667	694	686	654	**29**	611	514	421	338	240	**2,124**
Active Installed SATA Base Megawatt Power Draw		2,283	2,596	2,909	3,152	3,295	**30**	3,311	3,163	2,885	2,536	2,122	**14,016**
Total SAS Enterprise HDD K Units		3,111	1,890	765	0	0	**5,766**	0	0	0	0	0	**0**
SAS % ofTotal		3	2	1	0	0	**1**	0	0	0	0	0	**0**
Typical Operating Power Watts Per SAS Drive	8.6												
Typical Megawatt Power Draw for New Units		27	16	7	0	0	**50**	0	0	0	0	0	**0**
Active Installed SAS Base Megawatt Power Draw		308	226	165	88	45	**833**	23	0	0	0	0	**23**
Annual SSD Megawatt Power Draw		473	493	586	677	797	**3,026**	921	1,052	1,208	1,358	1,468	**6,007**
Annual HDD Megawatt Power Draw		618	684	701	686	686	**3,375**	611	514	421	338	338	**2,221**
Installed Base SSD Megawatt Power Draw		2,190	2,272	2,333	2,648	3,026	**12,470**	3,473	4,053	4,640	5,344	6,007	**23,516**
Installed Base HDD Megawatt Power Draw		2,591	2,823	3,075	3,240	3,340	**15,069**	3,334	3,163	2,885	2,536	2,122	**14,039**
Installed Base SSD + HDD Megawatt Power Draw		4,782	5,095	5,408	5,888	6,366	**27,538**	6,807	7,215	7,524	7,880	8,129	**37,555**
Installed Base Total "Active Archive" Megawatt Power Draw		25	24	23	25	27	**124**	29	31	35	37	40	**172**
HDD + SSD vs "Active Archive" Power Draw Ratio		191	212	235	236	236	**222**	235	233	215	213	203	**218**

Source: Furthur Market Research and Brad Johns Consulting (January 2024)

Notes relevant to Table 4 and Table 5:

- *A megawatt is one million watts. Watts-per-drive power draw has been converted to megawatts by multiplying typical per-drive operating power draw by millions of drives then dividing this sum by one million.*
- *The typical operating watts power draw per drive remains constant for NVMe, SAS, and SATA interfaces.*
- *For the active installed base, we assume a 5-year infrastructure refresh/replacement cycle, retiring all 2015 shipments in 2020 while adding 2020 shipments to the installed base of the prior year, and we repeat this cycle through 2035.*
- *The installed base per-drive watt power draw for SSDs and HDDs by interface has been determined by using the average interface percentage of the total drives shipped during the prior five years.*
- *Installed base "active archive" power draw is based on an average of 120 cartridges per drive and a maximum of 128 drives per library, combined with consistent integration of new, more power-efficient technologies, which results in lower power draw 2026-2030 and 2031-2035 compared with 2020-2025, despite enormous growth in the number of exabytes being managed.*

NVMe SSDs Draw 75% More Typical Operating Power Than SATA HDDs

It is a widely held belief that enterprise SSDs draw considerably less power than enterprise HDDs, but this is not the case on a per-drive basis. It is true that SATA and SAS SSDs draw less than 1/3 the power of SATA or SAS HDDs, but NVMe-interface enterprise SSDs draw more power per drive than SATA- or SAS-interface enterprise HDDs (11 watts as opposed to 6.3 or 8.6 watts, typical operating power).

Enterprise-grade SSDs and HDDs are in constant 24/7 power-on operation in enterprise data centers.[2] Rather than trying to determine with any degree of rigorous accuracy actual active/idle and read/write use patterns for SSDs and HDDs in diverse hot-warm-cool-cold-frozen data layers (an impossible task, and prone to great and variable error), we will simply use the "typical" watt operating power-draw specifications for NVMe, SAS, and SATA drives. While this analysis likely understates SSD power draw and overstates HDD power draw (since SSDs are utilized in far greater read/write/active operating than idle/standby mode in the hot and warm data layers, and HDDs are utilized in far greater idle/standby than read/write/active operating mode in the warm, cool, cold, and frozen data layers), it does yield meaningful comparative power-draw metrics from a constant benchmark for calibration.

Although power metrics per terabyte will decline for both HDDs and SSDs as average per-drive capacities expand, raw power draw watts per drive will not, especially when considering high-speed NVMe interfaces. Currently, NVMe interfaces are integrated in ~70% of all enterprise SSDs shipped and SATA interfaces are integrated in ~88% of all enterprise HDDs shipped. By 2029, we anticipate that NVMe SSD interfaces and SATA HDD interfaces will account for almost 100% of all enterprise drives shipped.[3]

New-unit power draw for SSDs actually exceeded the power draw for HDDs in 2021 and 2022. After the estimated drastic decline in enterprise SSD shipments 2023-2028—we currently believe these SSD markets will not recover to their 2022 unit-volume levels until 2029—HDD power draw will again exceed SSD power draw, but not by a wide margin. And the large delta in HDD vs SSD installed-base power draw diminishes greatly from 2021 onward. In 2020, the HDD installed base drew 2,492 megawatts, while the SSD installed base drew only 994 megawatts. In 2025, we estimate the HDD installed base will draw 2,501 megawatts, while the SSD installed base will draw 2,270 megawatts. In 2035, we estimate the HDD installed base will draw 2,122 megawatts, while the SSD installed base will draw 6,007 megawatts.

From 2020 through 2025, we estimate the combined SSD and HDD installed base will draw 25,267 megawatts. By staggering contrast, we estimate that the "active archive" installed base of tape and enterprise emerging storage infrastructures—including all rack, cooling, and robotic power requirements—will draw only 211 megawatts, less than 1% of the power required just to drive the bare bones SSDs and HDDs. We seriously doubt these power draw ratios can much change in coming years. In fact, this power draw ratio will increase in favor of tape and emerging enterprise storage infrastructures.

While SSD and HDD power draw per drive will remain roughly the same as their technologies evolve, in the tape and enterprise emerging storage infrastructures, the retirement of old technologies combined with the advent of more-efficient new technologies from 2023 through 2035 will require less power draw to service the exabytes forecast in Table 2.

As previously stated, at least 70%, more probably 80%, of all enterprise data is and will continue to be "cool" or "cold" or "frozen," with infrequent access times of minutes to days to weeks to years to decades, with little or no need for the performance of SSDs and HDDs, but with greatly expanding needs for Sustainability, Immutability, and Security (SIS), which SSDs and HDDs can neither cost effectively nor power efficiently fulfill.

There will be great growth in the tape and enterprise emerging storage infrastructures, but—without substantial change in purchasing and integration policies—it is still likely that huge numbers of HDDs and a significant number of SSDs will continue to manage far too many of the cool-cold-frozen workloads at far too great a cost per gigabyte while consuming an inordinate share of available energy.

---

2. *Cycling down banks of SSDs and HDDs to conserve power will result in mostly stillborn initiatives. Enterprise-grade SSDs and HDDs are designed for 24/7 power-on operation, and do not fare well in power-down mode or sitting on shelves. Powering up drives from power-down mode will increase failure rates and likely offset any cost savings derived from cycling down power conservation.*
3. *We doubt there will be any new interfaces for SSDs integrated in massive volume prior to 2035. We predict the effective death of the SAS interface for both SSDs and HDDs during 2028. Although much discussed and debated, we seriously doubt there will be any significant shipments of dual-actuator or NVMe-interface HDDs that would considerably increase HDD power draw metrics.*

V Sustainable Long-Term Data Management and Preservation

The Total Cost of Ownership (TCO) per TB of storing information in an active archive is substantially less today and will continue to have a compelling cost advantage in the coming years. Figure 8 depicts five-year $/TB trends from 2020-2035. Figure 9 depicts five-year kWh/TB installed base trends from 2023-2035.

Figure 8: Five-Year Costs-Per-Terabyte Trends, 2020-2035

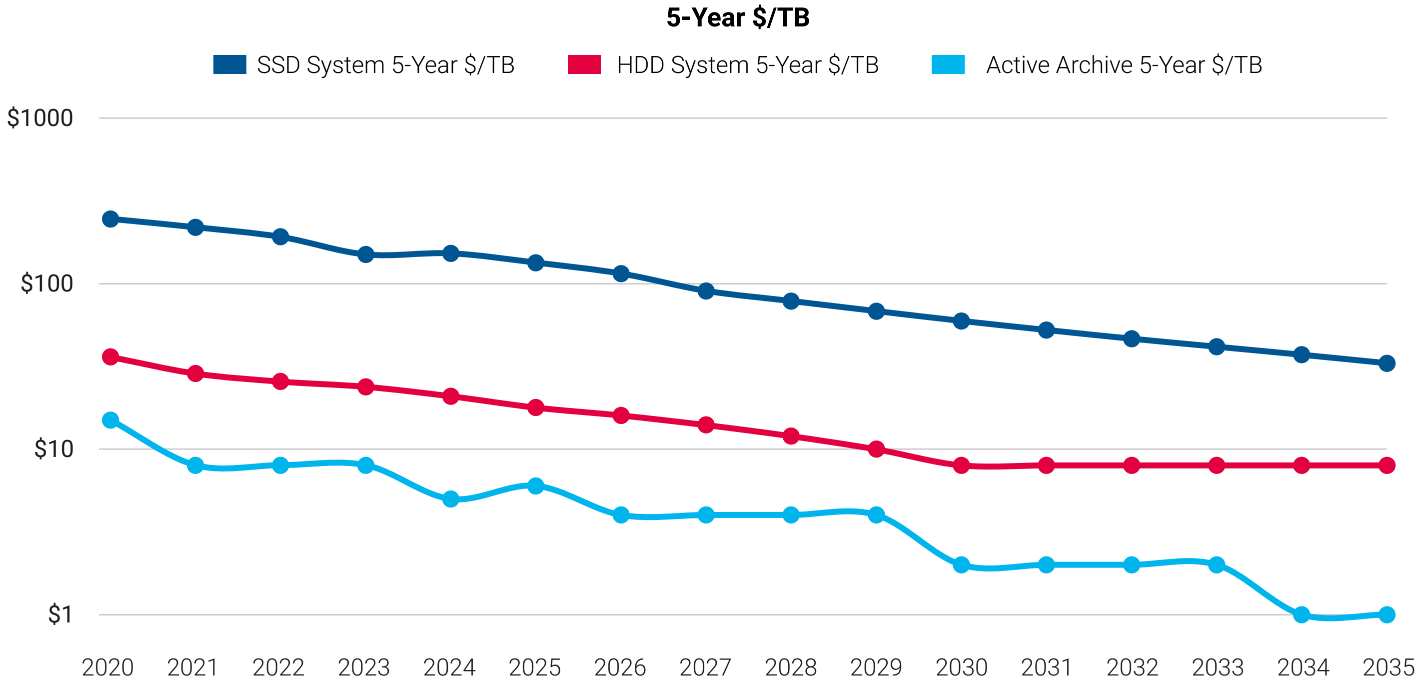

Note: Figure 8 uses a log scale. In 2020, 5-year costs per terabyte for an SSD system was 16.4x and for an HDD system was 2.4x the 5-year cost for an active archive system. In 2035, we project the 5-year costs per terabyte for an SSD system will be 33x and for an HDD system will be 8x the 5-year cost per terabyte for an active archive system.

Source: Brad Johns Consulting (January 2024)

Figure 9: Comparative SSD, HDD, and Active Archive System Installed Base kWh/TB Power Consumption, 2023-2035

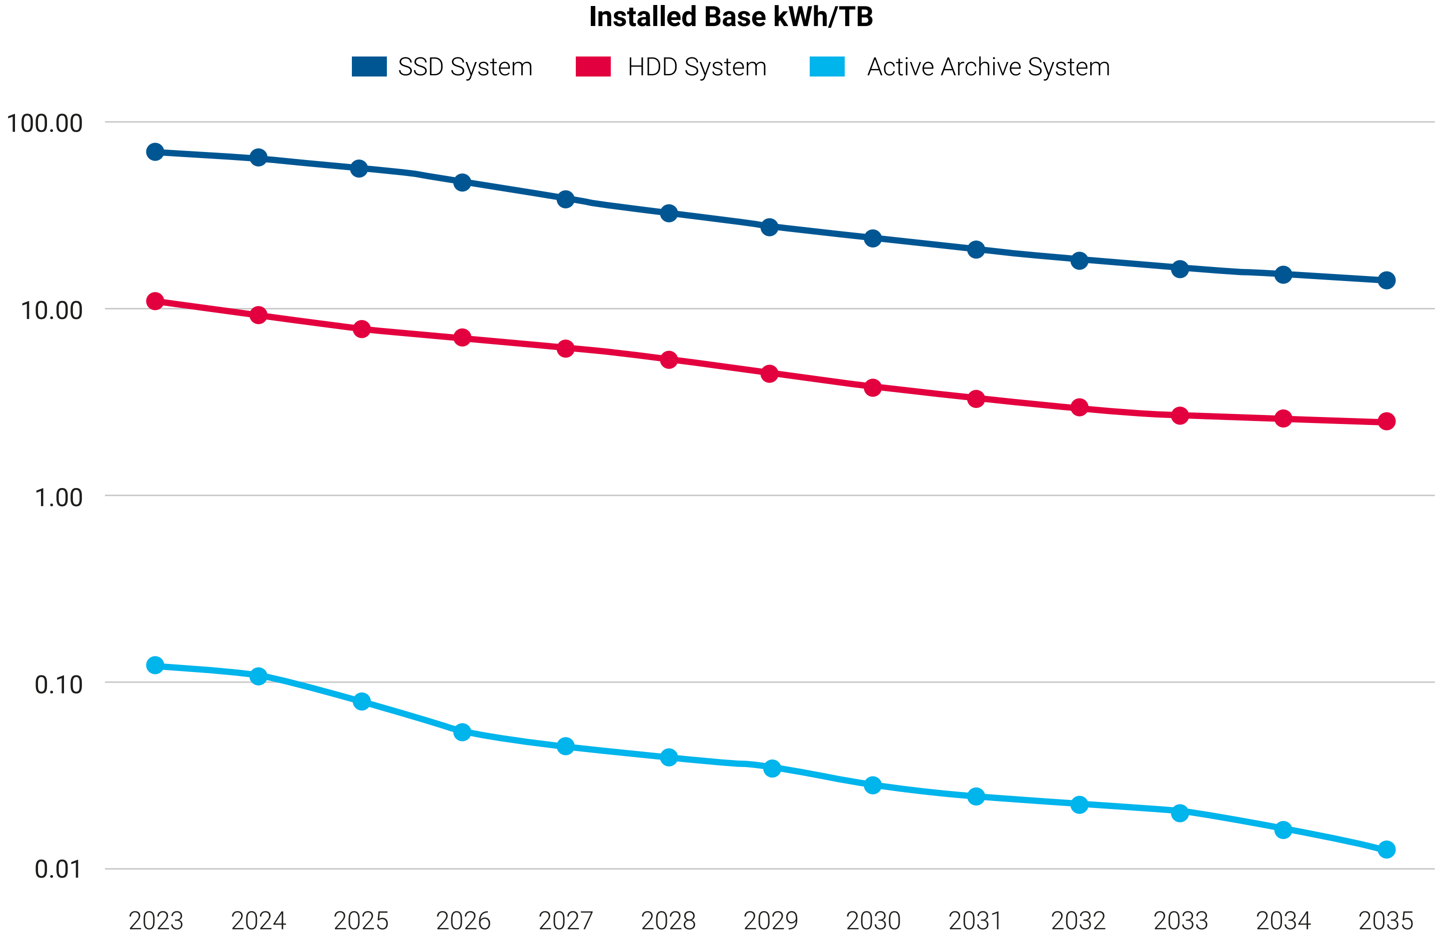

Note: Figure 9 uses a log scale. In 2023, the 5-year active installed base kWh/TB power consumption for an SSD system was 580x and for an HDD system was 90x that of an active archive system. In 2035, we project the 5-year installed base kWh/TB power consumption for an SSD system will be 1,000x and for an HDD system will be 200x that of an active archive system.

Source: Furthur Market Research and Brad Johns Consulting (January 2024)

Table 6 presents the end-user storage expenditures for each year based on the projected SSD, HDD, and active archive costs and capacities shipped. The costs include both the initial capital expense (CapEx) and operating expense (OpEx). OpEx includes the cost of the energy consumed and maintenance expense. All storage types, SSD, HDD, and active archive, are assumed to be refreshed every five years at the projected $/TB capital expense. This understates the cost of an SSD system, where a three-year life expectancy is standard, and overstates the cost of an active archive solution. We do not yet have any data for enterprise emerging technology systems, but tape drives are often kept for two LTO generations (around six or seven years) and tape libraries for over a decade.

Table 6 details our TCO analysis of SSD, HDD, and archive storage systems from 2020 through 2035, with projected costs of storing a terabyte of data in 5-year time periods.

The rapid growth of the dataverse creates not only energy consumption and CO2 emissions challenges but also cost challenges. There are several emerging technologies—namely, Cerabyte's ceramic nanolayers, Folio Photonics' dynamic multi-layer optical discs, Group47's DOTS (Digital Optical Technology System), Microsoft's silica, and DNA data storage—that collectively have the potential to exceed tape exabyte deliveries beginning in 2033. However, considering their immaturity, it is impossible to forecast their future cost and sustainability profiles. Therefore, as in section IV we are using our rigorous tape storage estimates as the basis to determine total "active archive" cost projections. Table 6 summarizes the costs associated with each technology including both the initial capital outlay (CapEX) and the annual operating expense (OpEx). Based on the projected industry shipments, the end user CapEx and OpEx for each storage type is projected out through 2035. In addition, for each year, we determine the $/TB cost of each storage type over a five-year period, consistent with the typical depreciation period for storage hardware.

Table 6: SSD, HDD, and Active Archive System TCO Estimates, 2020-2035

	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035
SSD System CapEx ($M)	$25,288	$30,861	$31,344	$14,095	$17,613	$20,964	$23,765	$23,505	$27,163	$30,225	$34,421	$38,111	$41,940	$45,905	$49,396	$51,515
SSD OpEx ($M)	$1,383	$1,673	$1,688	$1,080	$1,035	$1,119	$1,195	$1,219	$1,428	$1,614	$1,874	$2,144	$2,449	$2,811	$3,161	$3,417
SSD System 5-Year $/TB	$246	$219	$192	$150	$152	$134	$115	$90	$79	$68	$60	$52	$46	$42	$37	$33
HDD System CapEx ($M)	$16,780	$19,122	$16,846	$17,482	$19,928	$21,398	$24,265	$27,334	$28,394	$29,628	$29,516	$26,791	$22,558	$18,475	$14,873	$10,574
HDD OpEx ($M)	$1,546	$1,668	$1,454	$1,455	$1,690	$1,787	$2,079	$2,405	$2,455	$2,410	$2,257	$1,999	$1,684	$1,379	$1,110	$789
HDD System 5-Year $/TB	$36	$29	$26	$24	$21	$18	$16	$14	$12	$10	$8	$8	$8	$8	$8	$8
Active Archive CapEx ($M)	$1,772	$1,091	$1,174	$1,246	$1,085	$1,425	$1,544	$1,727	$2,516	$3,655	$3,038	$3,941	$4,950	$5,388	$4,569	$5,472
Active Archive OpEx ($M)	$59	$88	$93	$97	$70	$91	$77	$101	$151	$242	$160	$229	$302	$322	$234	$303
Active Archive 5-Year $/TB	$15	$8	$8	$8	$5	$6	$4	$4	$4	$4	$2	$2	$2	$2	$1	$1

Note: See Appendix for a much more-detailed analysis of these estimates.

Source: Brad Johns Consulting (January 2024)

The cost advantages of tape are clear and obvious. For example, the cost of storing a terabyte of data for a single five year (CapEx and OpEx) in 2023 on an SSD storage system is $150, on an HDD storage system is $24, and in a tape system is $8. Moreover, tape will have an enduring and widening cost advantage. Looking out to our 2035 estimates of future SSD, HDD, and tape shipments, the costs of storing a TB of data for five years in an SSD storage system will be $33, in an HDD system $8, and in a tape system $1.

Despite the projected increase in the use of active archive technologies in coming years, there is a significant opportunity available today to not only reduce power and carbon emissions but to dramatically reduce costs.

Using the 2023 CapEx and OpEx estimates, for every exabyte of cold or frozen data moved from HDD to tape storage, total costs can be reduced by more than $16 million over five years. The estimated annual energy costs will also drop by almost $1 million per exabyte, and reduce annual CO2 emissions by almost seven kilotons.

VI Inconclusive Conclusions

If the surging tide of stuff to be stored cannot be stemmed—and apparently it cannot—then new enterprise data infrastructures must not only cost less but must also consume less power to be in crucial and resilient alignment with the total availability of energy. This will require dramatic change in the purchasing and integration practices of myriad small and large data centers scattered throughout the world.

With the advent of new tape and enterprise emerging storage technologies, we have forecast that active archive shipments will expand to comprise more than 50% of the fresh enterprise zettabytes delivered in 2034 and 2035. These ~11 zettabytes will still fall far short of servicing the ~80% (~16 zettabytes) of new shipments in 2034 and 2035 that will be destined, within 60 days, to become cool or cold or frozen. And from 2031-2035, we project just the bare bones SSDs and HDDs in the installed base will draw ~218 times as much power as the installed base of active archive systems.

> **We fear there will continue to be immense waste of energy and money expended in the ways we choose to store and manage the active installed base of enterprise data. This will be tragic—"tragic" because the consequences of this waste can be so easily avoided.**

In the end, the CFOs, with fervent approval from the CEOs and board members, will have the final say. And in the cool and cold and frozen enterprise data layers—which have little or no real need for the performance of SSDs or HDDs, but have greatly expanding needs for Sustainability, Immutability, and Security—the most cost-effective and power-efficient technologies will inevitably prevail, because they make the greatest fiscal and ecological sense.

Appendix
More Detailed and Extensive TCO Estimates, 2020-2035

Appendix Table: SSD, HDD, and Active Archive System TCO Estimates, 2020-2035

	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035
SSD System CapEx ($M)	$25,288	$30,861	$31,344	$14,095	$17,613	$20,964	$23,765	$23,505	$27,163	$30,225	$34,421	$38,111	$41,940	$45,905	$49,396	$51,515
SSD System OpEx																
Energy Cost ($M)	$1,312	$1,587	$1,601	$1,024	$981	$1,061	$1,133	$1,156	$1,354	$1,530	$1,777	$2,033	$2,323	$2,665	$2,997	$3,240
Maintenance Cost ($M)	$72	$87	$87	$56	$54	$58	$62	$63	$74	$84	$97	$111	$127	$145	$164	$177
Total OpEx ($M)	$1,383	$1,673	$1,688	$1,080	$1,035	$1,119	$1,195	$1,219	$1,428	$1,614	$1,874	$2,144	$2,449	$2,811	$3,161	$3,417
SSD System Cost/TB ($)																
1-Year	$204	$182	$159	$117	$125	$111	$97	$76	$65	$57	$49	$43	$38	$34	$30	$26
3-Year	$225	$200	$175	$133	$139	$123	$106	$83	$72	$63	$54	$48	$42	$38	$34	$30
5-Year	$246	$219	$192	$150	$152	$134	$115	$90	$79	$68	$60	$52	$46	$42	$37	$33
10-Year	$380	$334	$282	$229	$221	$194	$168	$137	$120	$105	$93					
15-Year	$440															
HDD System CapEx ($M)	$16,780	$19,122	$16,846	$17,482	$19,928	$21,398	$24,265	$27,334	$28,394	$29,628	$29,516	$26,791	$22,558	$18,475	$14,873	$10,574
HDD System OpEx																
Energy Cost ($M)	$1,158	$1,250	$1,089	$1,090	$1,266	$1,339	$1,557	$1,801	$1,839	$1,805	$1,691	$1,498	$1,261	$1,033	$831	$591
Maintenance Cost ($M)	$388	$419	$365	$365	$424	$448	$522	$603	$616	$605	$566	$502	$422	$346	$279	$198
Total OpEx ($M)	$1,546	$1,668	$1,454	$1,455	$1,690	$1,787	$2,079	$2,405	$2,455	$2,410	$2,257	$1,999	$1,684	$1,379	$1,110	$789
HDD System Cost/TB ($)																
1-Year	$27	$22	$19	$18	$16	$14	$12	$10	$9	$7	$6	$6	$6	$6	$6	$6
3-Year	$32	$25	$23	$21	$18	$16	$14	$12	$10	$9	$7	$7	$7	$7	$7	$7
5-Year	$36	$29	$26	$24	$21	$18	$16	$14	$12	$10	$8	$8	$8	$8	$8	$8
10-Year	$54	$44	$40	$36	$31	$26	$24	$22	$20	$18	$16					
15-Year	$62															
Active Archive CapEx ($M)	$1,772	$1,091	$1,174	$1,246	$1,085	$1,425	$1,544	$1,727	$2,516	$3,655	$3,038	$3,941	$4,950	$5,388	$4,569	$5,472
Active Archive System OpEx																
Energy Cost ($M)	$1	$2	$2	$2	$2	$2	$2	$2	$3	$5	$4	$5	$7	$7	$5	$7
Maintenance Cost ($M)	$58	$86	$91	$95	$69	$89	$75	$98	$148	$237	$157	$224	$296	$314	$229	$296
Total OpEx ($M)	$59	$88	$93	$97	$70	$91	$77	$101	$151	$242	$160	$229	$302	$322	$234	$303
Active Archive System Cost/TB ($)																
1-Year	$13	$6	$6	$6	$4	$5	$3	$3	$3	$3	$2	$2	$2	$2	$1	$1
3-Year	$14	$7	$7	$7	$5	$5	$4	$3	$3	$4	$2	$2	$2	$2	$1	$1
5-Year	$15	$8	$8	$8	$5	$6	$4	$4	$4	$4	$2	$2	$2	$2	$1	$1
10-Year	$21	$12	$12	$11	$9	$8	$6	$6	$6	$5	$3					
15-Year	$23															

Source: Brad Johns Consulting (January 2024)

Notes:

- *End user storage expenditures are estimated for each year based on the projected SSD, HDD, and tape costs and capacities shipped. The yearly costs include both the initial capital expense (CapEx) and operating expense (OpEx). OpEx includes the cost of the energy consumed and maintenance expense. All storage types, SSD, HDD, and tape are assumed to be refreshed every five years at the projected $/TB capital expense. This understates the cost of an SSD system, where a three-year life expectancy is standard, and overstates the cost of a tape solution, wherein the tape drives are often kept for two LTO generations (around 6 or 7 years) and tape libraries for over a decade.*
- *The SSD and HDD system costs are based on industry standard rack configurations. The total system cost includes not only the projected costs of the SSD or HDDs, but also the cost of servers, JBODs, switches and power units. The capacity of the rack is increased by 20% for the SSD systems due to the use of compression in solid-state arrays. The capacity of the SSD and HDD rack is increased each year based on the projected average drive capacity.*
- *Tape system costs are based on publicly available pricing and estimates for a tape library (using the IBM TS4500 series), tape drives (LTO), and tape media. Future tape capacities are based on the LTO roadmap. LTO Generation 10 is projected to be available in 2025 with a new generation that offers double the capacity of the previous generation every three years throughout the analysis period. Future cartridge costs are projected to enter the market at $165 and decline by 20% every year to a floor of $65. Based on historical experience, tape drive costs are estimated to increase by 8% for each successive generation while tape library costs are constant. The product life cycle for tape drives is expected to be eight years, while tape cartridges have a longer life cycle of 12-15 years because organizations continue to purchase cartridges long after the manufacture of the generations of tape drives has stopped.*
- *For each generation of tape drives, a tape library (or libraries) for an exabyte of capacity is configured based on 120 tape cartridges/drive and 2.5 to 1 compression. For example, using LTO 7, one exabyte of storage requires four tape libraries, including four L frames, 32 D frames, 25 S frames, and 556 tape drives, along with 66,667 tape cartridges. In contrast, for the same amount of capacity, using LTO Generation 10, only one tape library is required, including 1 L frame, 6 D frames, 4 S frames, and 93 tape drives. The CapEx for the storage of 1 exabyte using LTO 7 is $11.7 million versus $3.7 million for LTO 10.*
- *For each year for each storage type, the $/TB of CapEx and OpEx is estimated over 1-, 3-, 5-, 10-, and if feasible 15-year time periods. For example, in 2020 the cost of storing 1 terabyte of data on an SSD system for one year is $204. This includes both CapEx and OpEx for one year. The cost of storing the 2020 data over three years is the cost of the original year plus two more years of OpEx. The cost of storing the 2020 data over five years is the cost of the initial year and four years of OpEx. The cost over 10 years includes the cost of the first five years, plus CapEx in year 6 to refresh the original devices, plus 5 years of OpEx. This is done for each storage type over the analysis period.*

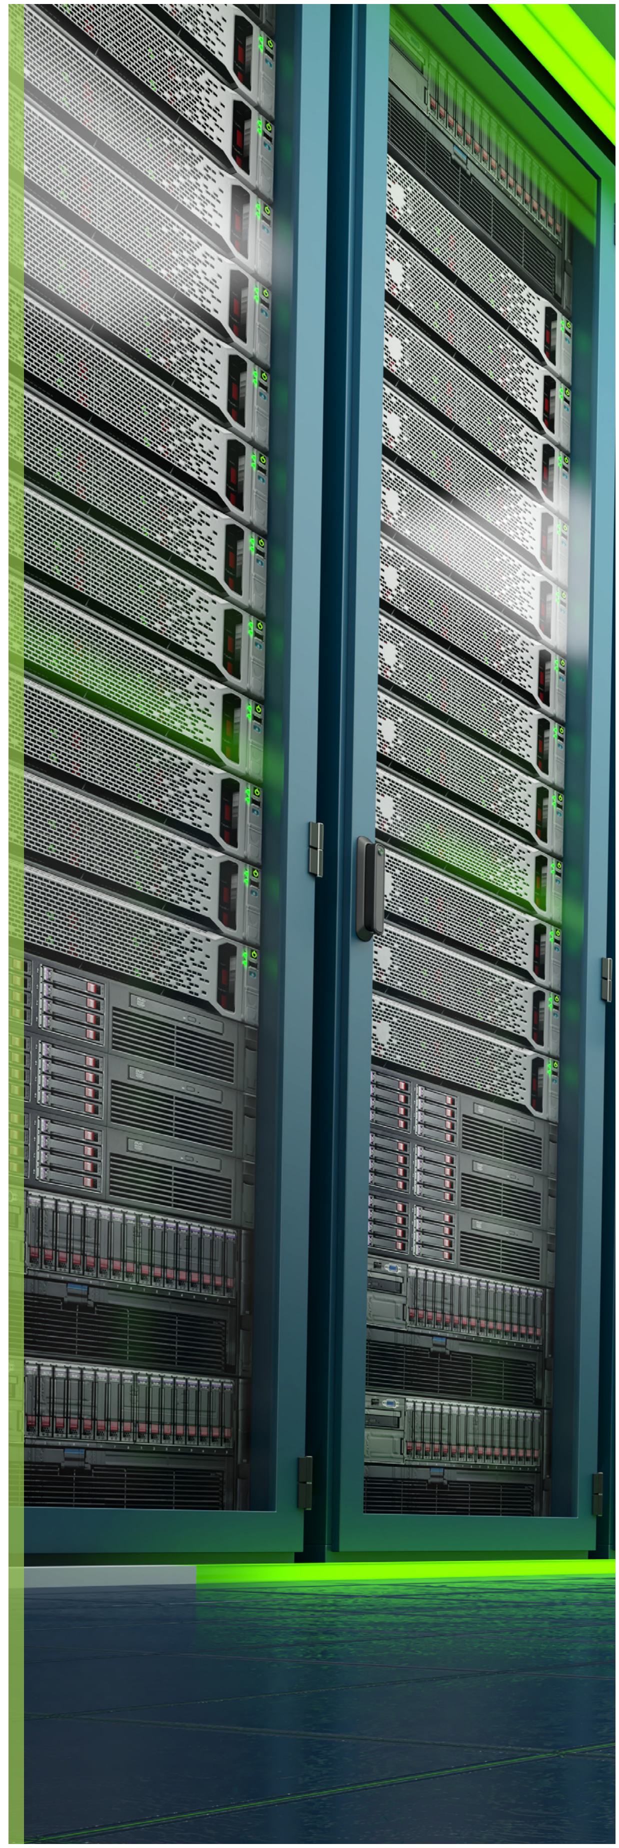

This whitepaper was jointly sponsored by Cerabyte, Fujifilm and IBM and was written by **John Monroe** and **Brad Johns.**

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

Brad Johns Consulting, LLC, **Email:** Brad.Johns@bradjohnsconsulting.com, **Website:** www.bradjohnsconsulting.com

Cerabyte is a deep-tech start-up developing a new ecosystem for sustainable low-cost long-term data storage at the verified prototype system level (TRL6). Ceramic-on-Glass media using laser-matrix writing and high-speed microscope reading technology is the foundation enabling the storage of vast amounts of cold data, that is kept for decades to centuries and rarely retrieved. Leveraging semiconductor tool R&D investment, Cerabyte is uniquely positioned to deliver the sustainable, accessible, and affordable EB-scale data center rack storage solutions required for the yottabyte era.

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.

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.

Monroe Biography

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

- From 1997 to 2022, Monroe was a VP Analyst at Gartner, covering the history and forecasting the future of consumer and enterprise storage markets.
- 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.
- 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).

Unlike most industry analysts, Monroe has had balance-sheet accountability for the stuff that he studies. 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.

Johns Biography

Brad Johns is President of Brad Johns Consulting LLC, an Information Technology consulting firm specializing in storage technology economics, marketing, and strategy. He has over 40 years of experience in the IT industry.

- In 2010, he established and led Brad Johns Consulting LLC, which provides consulting and marketing strategy assistance for storage technology companies.
- From 1997 until his retirement in 2010, he held various IBM storage product marketing and management leadership positions for IBM's worldwide disk, storage virtualization, tape storage, and archive product portfolio.
- From 1978 to 1997, after starting with the Data Processing Division, he held a variety of enterprise sales, sales management, industry marketing, and consulting positions. He was a founding member of IBM's Innovation Workflow consulting team and engaged with leading-edge clients in the aerospace, automotive, and information technology industries.
- Johns earned a BA in Economics from the University of Arizona in 1976 and a master's in business administration (MBA) in 1977.

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