

THE SUSTAINABLE PRESERVATION OF ENTERPRISE DATA



The Growing Immensity of a “Cold Yet Active Archive” (~80% of the installed base): ~35ZB in 2035...

Total Annual Enterprise Shipment Estimates

2023: 1.4ZB

2035: 10.3ZB

Total Active Installed Base of Enterprise Data

2023: 5.7ZB

2035: 43.3ZB

The billions of people and systems and sensors connected in the global dataverse have generated and will continue to generate immense quantities of data...~80% of which will become “cool-cold-frozen” within 60 days of their creation ...Our global digital culture’s obsession with and addiction to data could lead to a disproportionate increase in a power consumption metric...

*Whitepaper Extracts Presented by John Monroe at the Library of Congress DSA Conference
Monday 15 April 2024*

The complete whitepaper is available for free download at both of the websites cited below. Authors:

John Monroe, Furthur Market Research, LLC

torstones@gmail.com

Mobile Office: +1 (408) 828-3967

www.Furthurdata.com

Brad Johns, Brad Johns Consulting, LLC

Brad.Johns@bradjohnsconsulting.com


www.bradjohnsconsulting.com

OUR EXPANDING DATAVERSE



- *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...*

AGENDA



- Shipment History and New 2023-2035 Forecasts
- Updated Storage Pyramid
- A Few Power and TCO Estimates
- Inconclusive Conclusions

Enduring Question: *Will the Past be Prologue, or Will History Be Bunk?*

Note — My forecasts are always 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*

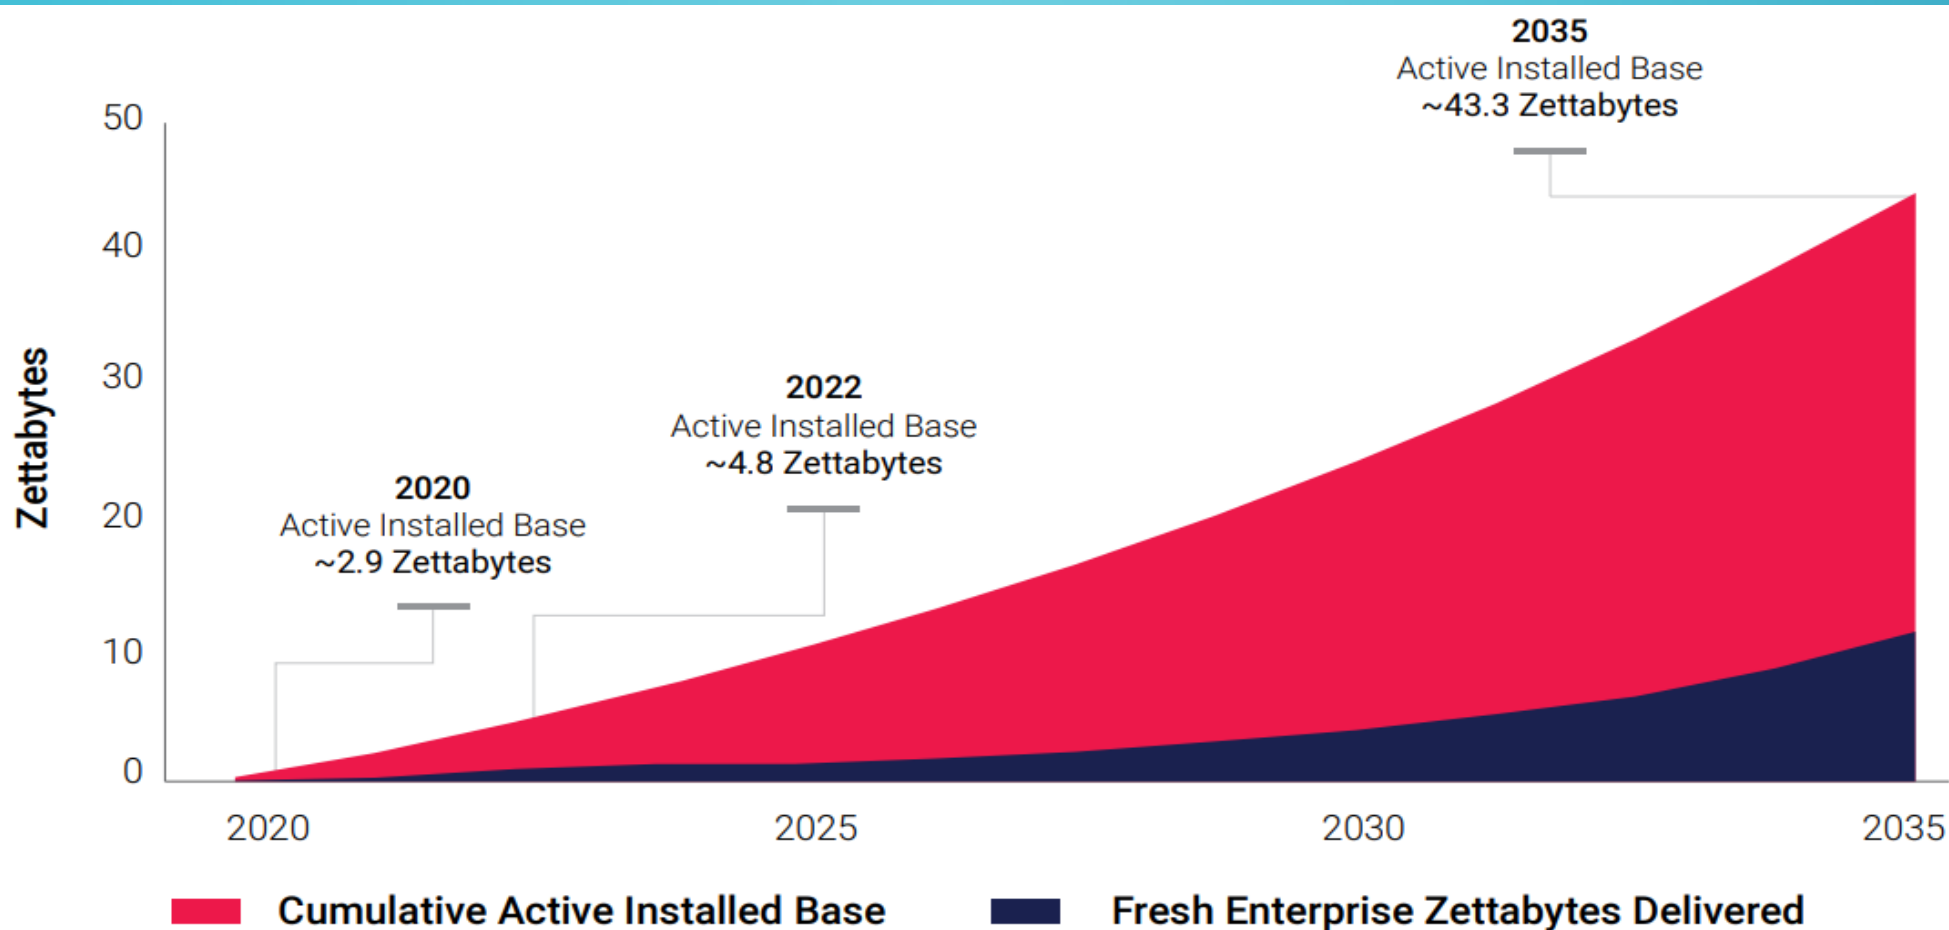
SHIPMENTS AND NEW FORECASTS, 2010-2035

	2010	2015	2020	2025	2030	2035	CAGR 2010-2022 (Actual)	CAGR 2023-2030 (Forecast)	CAGR 2031-2035 (Forecast)
Compressed* Enterprise SSD EB	0.2	13.5	130.8	198.2	735.6	2,074.7	47.8	28.1	22.2
Raw, Uncompressed Enterprise HDD EB	45.2	157.1	679.9	1,700.0	5,066.4	1,833.7	28.8	25.4	-20.7
Total SSD+HDD EB	45.4	170.6	810.7	1,898.3	5,802.0	3,908.4	30.9	25.7	-8.5
Compressed** LTO+IBM TS1100 "Active Archive" Tape EB	30.2	98.4	136.1	314.2	948.5	2,460.8	17.4	22.6	20.5
Enterprise Emerging Storage (2025 Onward)				18.8	735.0	3,904.8	-	108.3	40.6
Total "Active-Archive" Storage Opportunity				332.9	1,683.5	6,365.6	17.4	38.3	30.9
Total Compressed Shipments EB	75.6	269.0	946.8	2,231.2	7,485.5	10,274.0	27.2	27.1	7.3
SSD EB % of Total Annual Shipments	0.3	5.0	13.8	8.9	9.8	20.2			
HDD EB % of Total Annual Shipments	59.8	58.4	71.8	76.2	67.7	17.8			
Tape EB % of Total Annual Shipments	39.9	36.6	14.4	14.1	12.7	24.0			
Enterprise Emerging Storage EB % of Total Annual Shipments				0.8	9.8	38.0			
"Active-Archive" Storage % of Total EB				14.9	22.5	62.0			
Active Installed Base EB	91.0	819.9	2,923.2	8,039.3	24,924.9	43,310.1	39.2	23.5	9.9
Active Installed Base xChange Relative to 2010		9.0	32.1	88.3	273.9	475.9			

Source: Furthur Market Research (January 2024)


- 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 petabytes delivered in 2034 and 2035.
- But 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.

ENTERPRISE DATA SHIPMENTS AND THE ACTIVE INSTALLED BASE, 2020-2035



Source: Furthur Market Research (January 2024)

ALTERNATE ACTIVE INSTALLED BASE SCENARIOS, 2023-2035

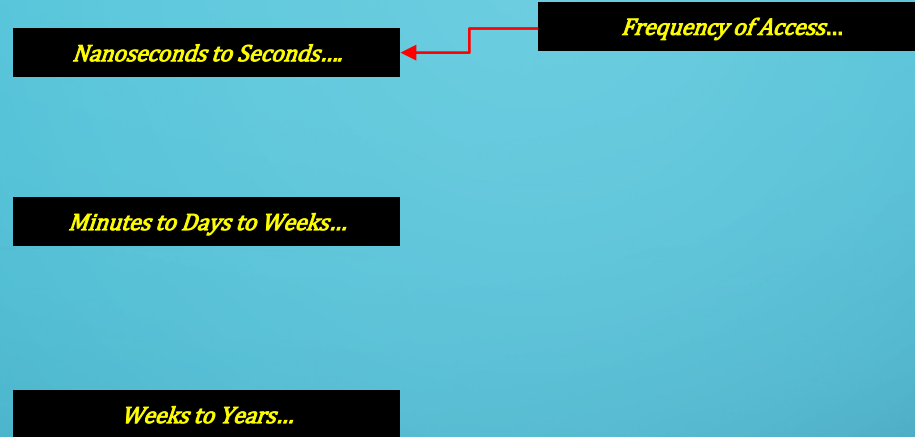


- Even in a worst-case 11%-per-year growth scenario, the active installed base in 2035 will have expanded to more than 20 zettabytes, more than 200x over 2010. In a 32% growth scenario—however unlikely, given recent results, but it should still be considered and not immediately dismissed—the active installed base might expand to more than 100 zettabytes in 2035.

THE EVOLVING STORAGE PYRAMID...

2023 Cool-Cold-Frozen Data: 75%
Shipments=1.05ZB
Installed Base=4.265ZB

2035 Cool-Cold-Frozen Data: 80%
Shipments=8.22ZB
Installed Base=34.65ZB



The growing enormity of "frostbitten" data...

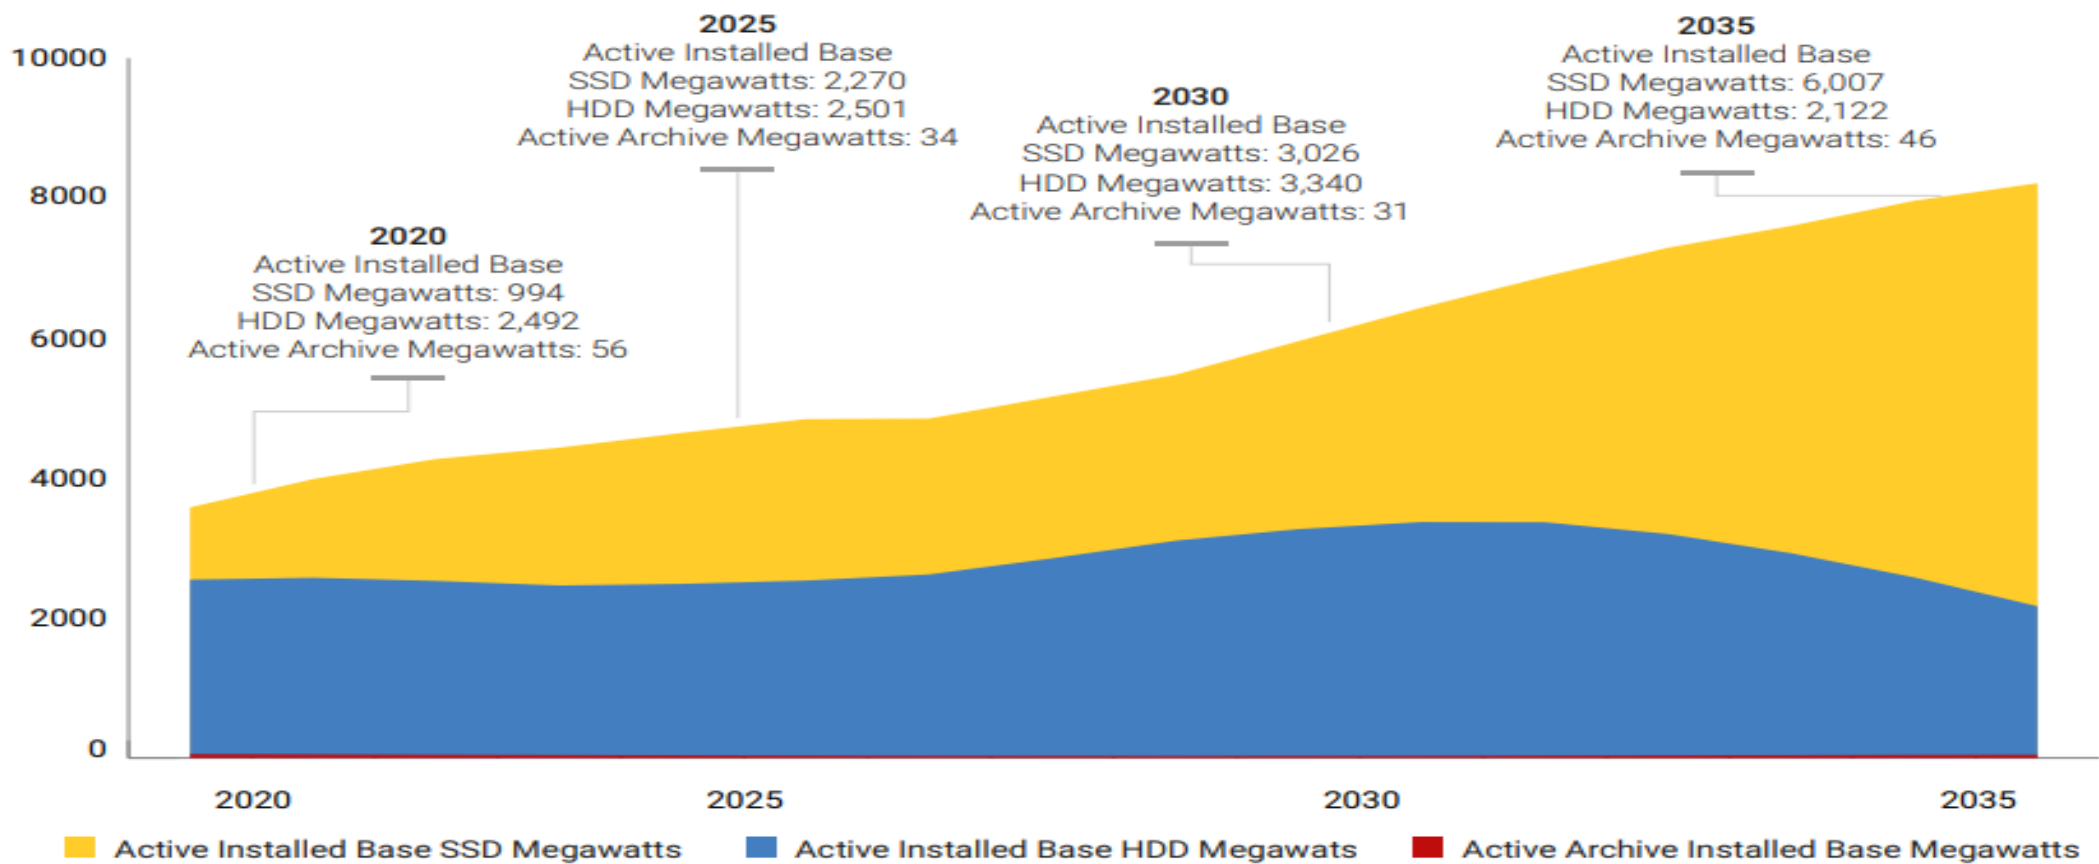
2023% of Total Exabytes Delivered:

SSDs 9.3% (2020:13.8%)
HDDs 74.4% (2020: 71.8%)
Tape 16.3% (2020: 14.4%)

Tape serviced only ~21.8% of the data destined to become, within 60 days, cool or cold or frozen...

**Σ *Blindingly blatant "fact":
Huge numbers of HDDs and a significant number of SSDs are managing and likely 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.***

POWER DRAW FOR BARE BONES SSDS AND HDDS COMPARED WITH ACTIVE ARCHIVE SYSTEMS, 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)

Despite great growth in tape and enterprising emerging storage technologies, 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*.

5-YEAR TCO TRENDS, COSTS/TB



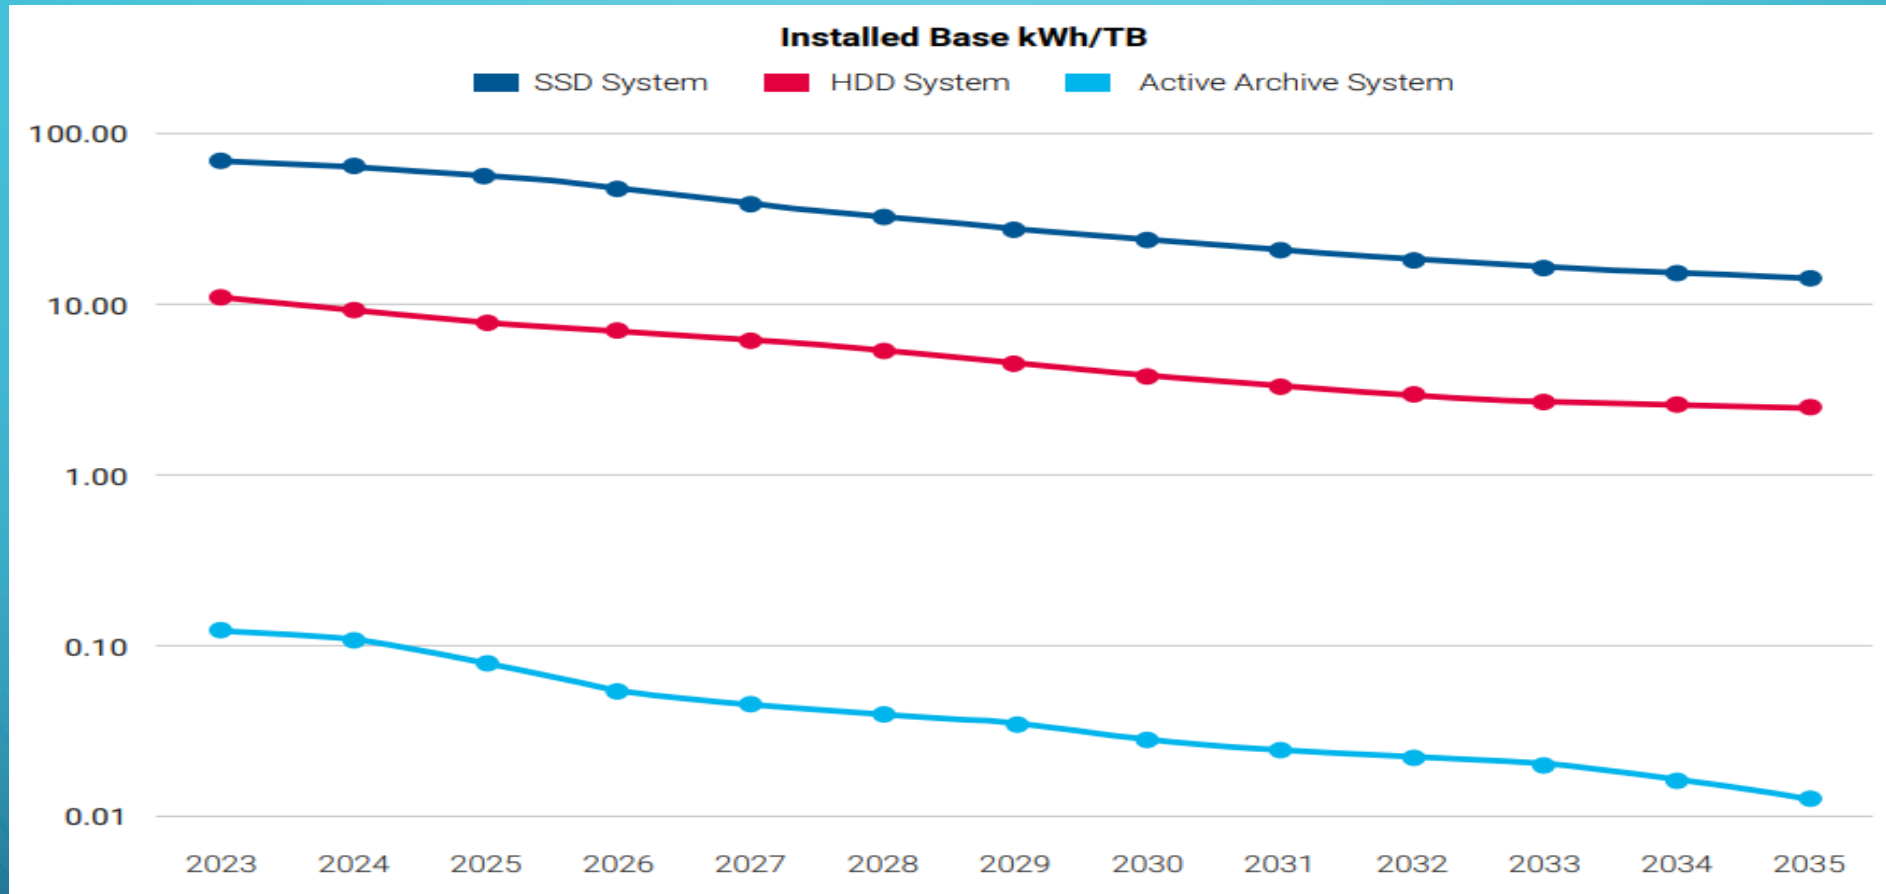
Source: Brad Johns Consulting (January 2024)

Note: This figure uses a log scale.

In 2020, 5-year costs per terabyte for an SSD system were $16.4x$ and for an HDD system were $2.4x$ 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.

5-YEAR KWH/TB INSTALLED BASE TRENDS, 2023-2035



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

Note: This figure 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.

INCONCLUSIVE CONCLUSIONS: EVOLVING DATA ECOSYSTEMS

- If the surging tide of stuff to be stored cannot be stemmed—and apparently it cannot, notwithstanding new uses of AI as a “gatekeeper”—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.
- 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 (SIS)—the most cost-effective and power-efficient technologies will inevitably prevail, because they make the greatest fiscal and ecological sense.
 - ❖ *One of my critics suggested I add “A” (access) or “SA” (seamless access) to “SIS”...SISA or SISSA...*
- In the end, the CFOs, with fervent approval from the CEOs and board members, will have the final say.




The background is a solid blue color with a subtle pattern of concentric, curved lines in a darker shade of blue. In the corners, there are white line-art graphics resembling circuit traces or data paths, with small circles at the end of the lines.

APPENDIX: A FEW NOTES TOWARD AN ESTIMATION OF SOME ACTUAL IMPACTS OF AI ON ENTERPRISE DATA

PRELIMINARY EXPECTORATION — IMPACTS OF AI/ML ON ENTERPRISE DATA: VASTLY INFLATED HYPE, ENORMOUS DANGERS, IMMENSE POSSIBILITIES

Before we work on
Artificial
Intelligence...

Why don't we do
something about
Natural Stupidity?



Accelerators


Endless Appetite for DATA: the more data that can be summoned and scrutinized, the better and more useful the AI/ML output

Obvious Vertical Markets: healthcare, surveillance, science, finance, manufacturing, retail...

Inhibitors

Gatekeeper: what does and does not need to be stored

Analysis: determining actual access needs in the hot-warm-cool-cold-frozen data infrastructures




We do not yet know what we do not know...

ARTIFICIAL INTELLIGENCE has the same relation to INTELLIGENCE as artificial flowers have to flowers




A year spent in artificial intelligence is enough to make one believe in God.

— Alan Perlis —



A DEEPLY HOPED-FOR EVENTUALITY: ARTFUL SYMBIOSIS ON A GLOBAL SCALE

"NATURE DOES NOT KNOW EXTINCTION, ALL IT KNOWS IS TRANSFORMATION"...



- *Computer graphics are no more generated by a computer than the Sistine Chapel ceiling was created by brushes and pigments. But will AI change this game, is it really different this time?*
- *AI: Like any applied science, once it's seen to work it can't be stopped...but can it be regulated, controlled?*
- *Inevitably and inescapably, AI will reveal, once again, man's immense capacity for evil. But hopefully we will also discover, once again, that there is more in man to admire than there is to despise.*
- *The future uses of storage technologies: An enlargement of the library of forms in which DATA, unleashed in fresh dimensions, can come to profitable life...*