Improving Data Center Sustainability with Modern Tape Storage

Brad Johns
brad.johns@bradjohnsconsulting.com
August 2025

Brad Johns Consulting LLC

Executive Overview

The explosive growth of Generative Artificial Intelligence (GenAI) is reshaping digital infrastructure demands, accelerating data center construction, and driving energy consumption to unprecedented levels. With daily GenAI users numbering over 100 million, hyperscale data centers are deploying high-performance Graphics Processing Units (GPUs) and flash solid-state drive (SSD) storage at scale, resulting in data center power consumption that could rise to as much as 12% of total U.S. energy use by 2028[1].

Amid this surge, inactive or infrequently accessed data—estimated to comprise 60% to 80% of enterprise storage—poses a significant challenge to sustainability and cost. Storing cold data on hard disk drives (HDDs) substantially increases power consumption, carbon emissions, and electronic waste at a time when corporate sustainability commitments are under mounting scrutiny from stakeholders and regulators.

This report demonstrates that modern tape storage offers a compelling opportunity to reduce the total cost of ownership and dramatically cut emissions. A modeled scenario shows that moving inactive data to tape can lower associated CO_2 emissions by up to 95% compared to HDD-only storage. Active-archive strategies using a hybrid of tape and disk can still achieve considerable, double-digit reductions in emissions and costs.

As enterprises adapt to the demands of GenAI and intensifying energy compliance regulations, transitioning inactive data to modern tape storage provides a proven, scalable solution to control costs, reduce environmental impact, and support long-term sustainability goals.

Generative AI Has Dramatically Impacted the Digital Infrastructure

With the release of ChatGPT 2.5, the world's digital infrastructure has been set on a new trajectory. GenAI is making unprecedented demands on data center servers, storage, networking, cooling, and power consumption. ChatGPT is the fastest-growing consumer application in history. It is estimated that the number of ChatGPT users grew to over 100 million in just 64 days. Its growth has continued unabated, reaching over 400 million weekly active users in April 2025. Other GenAI tools have also entered the market and experienced dramatic growth. Google's Gemini is estimated to have over 47 million active users[2] through the first half of 2025. Microsoft CoPilot has been deeply integrated into the Office Suite of tools, making comparisons difficult, but estimates include 28 million active users as of October 2024. Worldwide, it is estimated that there are 115 million to 180 million individuals who use GenAI daily[3].

The development and deployment of GenAI applications make unparalleled demands on the IT infrastructure. Model development requires the use of very high-performance GPUs, flash memory, and high-speed networking. Massive amounts of data are input into the model, often petabytes of raw video, image, and text.

This data is then iteratively defined to develop a Large Language Model (LLM). All the inputs and iterations are frequently saved, creating vast amounts of information that must be archived for potential future access. In addition, once the model is complete, it is widely deployed to the edge to minimize network traffic.

To accommodate these demands, hyperscale data centers are integrating GPUs and high-performance storage, and many new data centers are being built specifically to support the development and deployment of GenAI. Data-center capacity is often measured by the amount of power consumed. Historically, data centers required 20 to 30 Megawatts (MWs) of power, and a typical server rack requires 7kW of power. With GenAI, data center power can easily be ten times greater, and racks supporting GPUs may require 30kW per rack.

GenAI has driven unprecedented growth in data center power consumption. According to the Lawrence Berkeley National Laboratory (LBL) report on United States Data Center Energy Usage[4], U.S. data center energy consumption grew from 76 TWh, approximately 1.9% of U.S. energy consumption in 2018, to 176 TWh, representing 4.4% of U.S. energy consumption in 2023. Furthermore, as GenAI continues to grow, it is projected that data center energy consumption could increase to between 325 and 580 TWh by 2028. This would represent 6.7% to 12.0% of the U.S.'s total energy consumption.

The amount of power consumed by storage is also projected to grow dramatically. Two storage types are included in the LBL study: hard disk drives and flash storage. In 2014, it was estimated that storage consumed 7 TWH of power, with HDDs comprising 96% of the total. Power consumption more than doubled by 2023 to 16 TWh, and Flash Storage accounted for 25% of the total. In 2028, it is projected that storage power consumption

---

[1] Berkeley Lab – 2024 United States Data Center Energy Usage Report[2] Furthur Market Research - The impacts of Generative AI on Enterprise Data, https://furthurdata.com/

[2] BusinessofApps - https://www.businessofapps.com/data/google-gemini-statistics/

[3] Technollama.co.uk - https://www.technollama.co.uk/a-gemini-report-how-many-people-are-using-generative-ai-on-a-daily-basis-a-gemini-report

[4] Lawrence Berkeley National Lab - 2024 United States Data Center Energy Usage Report- https://eta-publications.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report

will increase to 22 TWh, with flash accounting for 11 TWh of the total. The Flash storage demand is driven by the rapid expansion of GenAI and the need for high-performance computing to develop and utilize large language models.

Artificial Intelligence is also driving dramatic growth in industry storage shipments. From our latest report[5], we estimate that enterprise storage shipments for all storage media—SSD, HDD, Tape, and Emerging Technologies—will increase from 2,024 Exabytes (EB) in 2025 to 6,155 EB in 2030, a 24.9% CAGR. This results in the active installed base of enterprise storage growing from 7,770 EB in 2025 to 20,219 EB in 2030, a 21.1 % CAGR. The remarkable growth in the active installed base will raise complex sustainability challenges, not only in terms of power consumption, but also in terms of embedded carbon emissions and electronic waste (e-waste), at a time when corporate leaders are increasingly concerned about their companies' sustainability.

Industry Focuses on Carbon Emissions and Energy Consumption

Corporations are increasingly viewing sustainability as a key business initiative despite the world's economic and political turmoil. A recent Forbes Research survey of over 1,000 executives found that 28% of the leaders considered sustainability their top corporate agenda, up from 0% three years ago, and 39% more considered it among the top three items, up from 22% three years ago. In addition, American citizens are also concerned about climate change and believe that big business is not doing enough to combat it. A recent Pew Research study[6] found that 69% of Americans felt that companies were doing too little to combat climate change.

The Information Technology (IT) industry has responded to the demand for improved sustainability. Microsoft has committed to robust programs[7] to reduce its Scope 1 and 2 emissions (direct and indirect from energy consumption) to near zero and its Scope 3 (indirect) emissions by more than half by 2030.[8] IBM has committed to producing 75% of the electricity it consumes worldwide from renewable sources by the end of 2025 and 90% by 2030. Additionally, IBM aims to reduce its operational greenhouse gas (GHG) emissions by 65% by the end of 2025 compared to the base year of 2010 and to achieve net-zero operational GHG emissions by 2030. Seagate, a major HDD supplier, has announced two significant goals: to power its manufacturing and R&D sites with renewable energy by 2030 and to achieve carbon neutrality by 2040.[9] Amazon has committed to reducing its environmental impact and continues to make progress toward attaining net-zero carbon by 2040, and water positivity by 2030.[10]

Massive Amounts of Stored Data are Rarely Accessed

As noted above, a massive amount of digital information is being created, transmitted, replicated, stored, and managed. Yet, ironically, despite the enormous growth of information, much of the information is "cold data," which is rarely accessed but still has value and cannot be deleted. In our January 2024 Report, "The Sustainable Preservation of Enterprise Data,"[11] we estimated that as much as 60%-80% of all stored information can be classified as cool, cold, or frozen, with access frequencies ranging from minutes to days, weeks, or years. Yet, the data often remains on HDDs, consuming an inordinate share of available power.

Storing cold data on tape rather than HDDs represents a significant opportunity for data centers to reduce their energy consumption, carbon emissions, and e-waste, while lowering the Total Cost of Ownership (TCO). To illustrate this opportunity, the environmental impact of a hypothetical data center that stores 100 PB of data with a ten-year retention requirement is calculated. First, the CO2e emissions resulting from keeping all the data on HDDs are estimated. Second, the CO2e emissions of an active-archive solution where 60% of the data is placed on tape and 40% on HDDs are calculated. Finally, the CO2e emissions for a deep-archive solution where 100% of the information is placed on tape are estimated. The analysis includes the CO2e generated due to energy consumption during the media's use, as well as the emissions associated with the acquisition of raw materials and the manufacturing of the storage media. A note of caution: this analysis includes only the emissions associated with the storage media and read/write mechanism, and not the supporting IT infrastructure, such as controllers, libraries, servers, networks, and supporting infrastructure.

In this hypothetical example, storing 100% of the data on hard disk drives generates 1,111 Metric Tons (t) of C02e

---

[5] Monroe and Johns, Further Market Research - ALL DATA IS INDISPENSABLE: The Staggering Immensity of an Active Archive - https://furthurdata.com/

[6] Pew Research - https://www.pewresearch.org/science/2024/12/09/how-americans-view-climate-change-and-policies-to-address-the-issue/

[7] https://www.microsoft.com/en-us/corporate-responsibility/sustainability/carbon-reduction

[8] https://newsroom.ibm.com/2021-02-16-IBM-Commits-To-Net-Zero-Greenhouse-Gas-Emissions-By-2030

[9] https://www.seagate.com/esg/planet/

[10] https://sustainability.aboutamazon.com/2024-amazon-sustainability-report.pdf

[11] Monroe and Johns, Brad Johns Consulting - https://www.bradjohnsconsulting.com/publications

over ten years. An active archive solution that stores 60% of the information on tape while retaining the active 40% of the data on HDD storage generates 480 t of C02e, a 57% reduction. A deep-archive solution that stores all the data on tape generates only 60 t of CO2e, a 95% reduction.

Figure 1 below illustrates the savings.

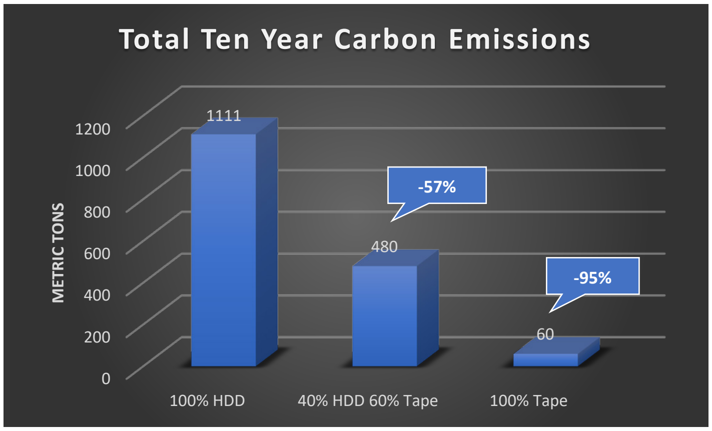

FIGURE 1 - TEN-YEAR CO2e EMISSIONS

In addition to lower carbon emissions, moving data to tape substantially reduces the amount of electronic waste (e-waste). A ten-year retention requirement will require at least one HDD refresh, resulting in the purchase of replacement HDDs in year 6. Based on HDD historical capacity growth, the year-six refresh is estimated to utilize 50TB HDDs. The larger-capacity HDD reduces the number of HDDs required, as well as the resulting CO2e emissions and electronic waste. Storing all 100 PB of data on hard disk drives and refreshing them after five years generates 4.6 t of e-waste. Tape storage has a longer life and does not require refreshment for over ten years. As a result, storing 60% of the data on tape results in only 2.3 tons of electronic waste, a 50% reduction in waste. Keeping all the data on tape produces only .75 t of e-waste, an 84% reduction.

Figure 2 illustrates the waste reduction.

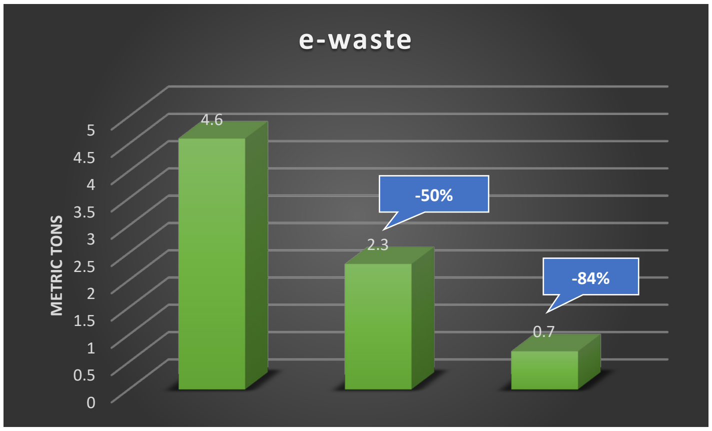

FIGURE 2 - TONS OF eWASTE

Moving inactive data to tape storage also dramatically reduces energy consumption. Referring to Figure 4, ten-year energy consumption from a 100% HDD solution is 2484 Megawatts (MW), for an active archive composed of 40% HDD and 60% tape, energy consumption drops to 1046 MW, a 58% reduction, and for a 100% deep-archive tape solution, power consumption drops to 88 MW, a 96% reduction.

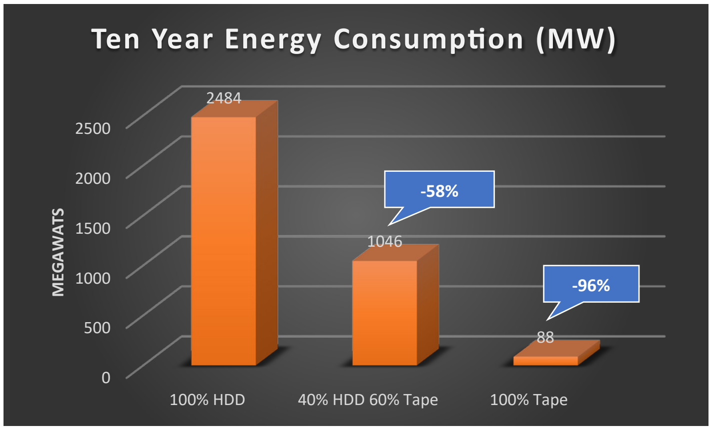

FIGURE 3 - POWER CONSUMPTION

Total Cost of Ownership

While the sustainability benefits of storing cold data on tape are substantial, there is also a significant financial benefit to be gained. To estimate the economic benefit, using the recently updated Fujifilm TCO tool, the TCO for a disk storage solution utilizing high-capacity HDDs is compared to that of a tape solution using LTO 10 tape media and drives. The estimated ten-year TCO for keeping all 100 PB on disk is $12,178,542. With the tiered solution, the ten-year TCO drops to $7,658,188, a 37% savings. Storing all 100 PB of data on tape reduces the TCO even further, decreasing the cost to $4,520,062, a 63% reduction.

Figure 4 illustrates the savings.

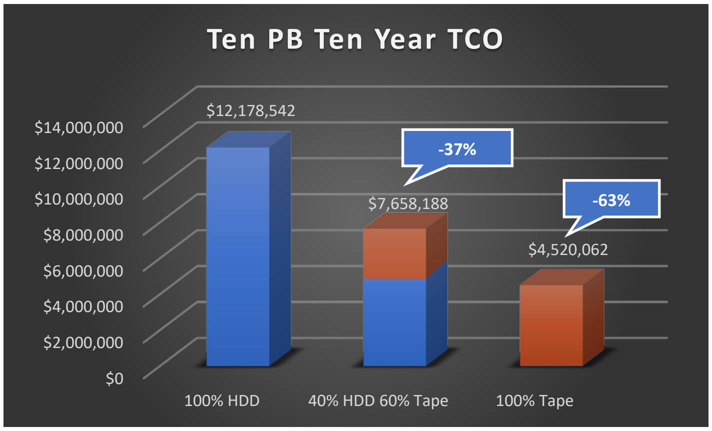

FIGURE 4 - TEN-YEAR TCO

Potential Global Impact

Potential sustainability and financial benefits are dramatic for a single firm. But what if organizations worldwide assessed their operations and migrated cold data to modern tape storage? We can develop an estimate of the potential worldwide impact by starting with the total amount of data stored globally and the media on which it resides. Collaborating with John Monroe of Further Market Research in our recent report[12], we estimated the global enterprise HDD active installed base would expand to 5.39 zettabytes (ZB) by the end of 2025.

What if a significant percentage of the disk-resident data was cold data and moved to tape storage?

Figure 5 illustrates the worldwide 10-year CO2e emissions of varying percentages of HDD-resident data being transferred to tape.

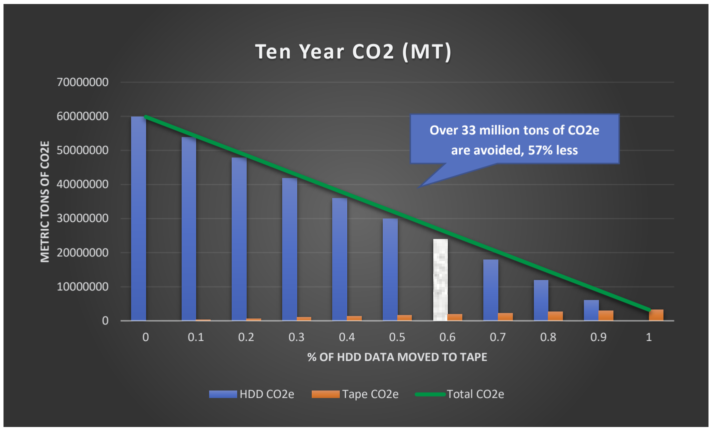

FIGURE 5 - TEN YEARS CO2E WITH % HDD DATA MOVED TO TAPE

For example, migrating 60% of the HDD-resident data to tape reduces CO2e emissions by over 33 million metric tons, a 57% savings.

Summary

The rise of GenAI is reshaping enterprise IT infrastructures. With an estimated 115 to 180 million daily users, GenAI is fueling rapid growth in data creation, driving the construction of hyperscale data centers, and pushing energy consumption to unprecedented levels. The Lawrence Berkeley National Laboratory projects U.S. data center energy use could more than triple by 2028, consuming as much as 12% of total national electricity. A significant portion of this challenge stems not from active workloads, but from the enormous volume of inactive or infrequently accessed data. Studies show that 60%–80% of enterprise data is "cold," yet it is often stored on HDDs, which contributes to increased power demand, carbon emissions, and e-waste waste.

This report demonstrates that modern tape storage offers IT leaders a practical and impactful path forward. Tape is purpose-built for inactive data: it consumes no energy when idle, generates far less embedded carbon, and requires fewer refresh cycles than HDDs. In modeled scenarios, shifting all inactive data to tape reduces CO_2 emissions by up to 95% compared with HDD-only solutions. Even active-archive strategies—where 60% of data is placed on tape and 40% remains on disk—achieve a 57% reduction in CO2 emissions and a 50% reduction in e-waste.

The financial advantages are equally compelling. Using the Fujifilm TCO model, a 100-PB, 10-year storage scenario reveals that disk-only solutions incur costs exceeding $12 million. A hybrid active archive lowers TCO by 37%, while a full tape deployment cuts costs by 63%.

As enterprises expand their GenAI capabilities and face intensifying scrutiny of sustainability commitments, the economic and environmental case for tape is clear. Migrating cold data to modern tape storage reduces risk, supports compliance with environmental, social, and governance (ESG) goals, and provides IT storage managers with a scalable, proven solution that delivers both cost control and measurable climate impact.

[12] Monroe and Johns, Further Market Research - ALL DATA IS INDISPENSABLE: The Staggering Immensity of an Active Archive - https://furthurdata.com/

Appendix

Notes on This Update

This report revisits the sustainability analysis published in the summer of 2022 and incorporates the new LTO Generation 10 and the higher capacity Seagate 30TB Mozaic 3+ drive. LTO tape CO2e/TB emissions improved primarily because the LTO 10 tape cartridge's CO2e profile is similar to that of the LTO 9 tape cartridge, with 7.94 kg of embedded carbon versus 7.63 kg for LTO 9. However, with a 30 TB capacity for LTO 10 versus 18 TB for LTO 9, the CO2e/TB is significantly lower. Additionally, the cartridge weight is identical, resulting in fewer tape cartridges for a given capacity requirement and reduced e-waste.

IBM provided LTO 10 tape drive embedded carbon estimates, which have been incorporated into the CO2e estimates. The embedded carbon for an LTO 10 tape drive is 48 kg. We assumed that one tape drive was required for every 120 tape cartridges.

The HDD information is updated to include the new 30 TB Seagate Mozaic 3+ HDD. The estimates of embedded carbon and power consumption were obtained from the Seagate "Decarbonizing Data" report[13]. For the CO2 and TCO estimates, we assume Erasure Coding with 80% efficiency was utilized for storing the data on HDDs.

CO2e Emissions

Fortunately, the data storage industry provides excellent data on the environmental impact of manufacturing, using, and disposing of its products. As mentioned above, the basis for the hard disk drive CO2e is the Seagate 30TB Moziac 3+ HDD. Given historical HDD capacity growth rates, it is projected that all the initial HDDs will be replaced in year 6 with 50 TB capacity HDDs with the same environmental footprint as the existing device, resulting in a significant reduction in the CO2e/PB.

Fujifilm provided estimates for the lifecycle CO2e emissions of LTO 9 tape media. For a comprehensive picture of carbon emissions. IBM provided the embedded carbon and energy consumption for the LTO 10 Full Height Tape Drive.

CO2e - kg

Bill of Materials	3.49
Manufacturing Energy	4.21
Packaging	0.24
Total	**7.94**

FIGURE 6 - LTO 10 TAPE MEDIA EMBEDDED CO2e

[13] Seagate – "Decarbonizing Data" – April 2025 - https://www.seagate.com/resources/decarbonizing-data-report/

Regarding the Information Used in This Report

Brad Johns Consulting LLC believes that the information in this report was accurate as of the date of publication. Information is provided "AS IS" without warranty of any kind.

Trademarks and Special Notices

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HPE, IBM Corp., and Quantum in the U.S. and other countries.

Other companies, products, or service names may be trademarks or service marks of others.

About the Sponsor

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.

About the Author

Brad Johns is the owner and President of Brad Johns Consulting LLC. He has over forty years of experience in the Information Technology industry. His firm specializes in economic analysis and consulting for the storage industry. He began his information technology career with the Data Processing Division of IBM in 1978, holding various sales management, consulting, and marketing positions. He joined the IBM Storage Systems Division in 1997 and was responsible for product management and marketing until his retirement in 2010. He holds a Master of Business Administration and a Bachelor of Arts degree in Economics from the University of Arizona.

Below is his contact information.

Email: brad.johns@bradjohnsconsulting.com
Website: bradjohnsconsulting.com

This paper was sponsored by IBM