Brad Johns Consulting, LLC

Improving Information Technology Sustainability with Modern Tape Storage —2022 Update

> Brad Johns brad.johns@bradjohnsconsulting.com August 2022

Executive Summary

As the world has made significant progress against the COVID pandemic and world economies rebound, climate change and the broader subject of Environmental, Social, and Governance (ESG) have become critical issues for companies worldwide. Nowhere is this experienced more than in the Information Technology (IT) industry, given its size and impact on consumers worldwide. The IT industry is rapidly changing with the evolution of cloud computing, artificial intelligence, big data, and the expansion of 5G networks resulting in the creation of massive quantities of information. Industry analysts estimate that the amount of data stored may grow to 17 zettabytes (ZB) by 2025¹. Most of this information ultimately resides on disk or tape storage systems. Given the focus on sustainability and the large volumes of storage devices required to store the growing quantities of data in the coming years, organizations have an opportunity to reduce their carbon footprint, improve sustainability and reduce expenses by migrating less frequently accessed (cold data) from hard disk drive (HDD) based storage to modern tape storage. Industry analysts believe that 60-80% of the data stored on HDDs is cold data. Consider a hypothetical example, assessing just the impact of the storage media of keeping 100 PB of information for ten years. Compared to an HDD-only solution, an active archive that moved 60% of the HDD resident data to tape reduced carbon emissions by 58%, reduced electronic waste by 51% and

reduced Total Cost of Ownership (TCO) by 46%. In addition, if all the data is cold and is transferred to tape, carbon emissions are reduced by 97%, electronic waste decreased by 85%, and TCO lowered by 78%. Globally, moving 60% of the HDD resident data to tape could reduce global carbon emissions by 79 million tons, making a meaningful contribution to reducing global carbon emissions. For enterprises, the financial benefits are significant enough to raise the possibility of funding the migration project entirely by the TCO savings while helping reduce global carbon emissions and electronic waste.

Renewed focus on climate change

Climate change and sustainability have re-emerged as significant issues. In 2020, the COVID pandemic, the global recession, and the highly contested elections in the United States grabbed the headlines. With worldwide COVID vaccinations providing light at the end of the pandemic tunnel, world economies beginning to recover, and a change of administration in the United States, the importance of combating climate change has once again moved front and center.



The financial community has also taken note of the potential impact of climate change. The United States Securities and Exchange Commission issued a request for public comment on climate disclosures². The SEC initiative seeks to capture ideas from all constituents to formulate a standard set of disclosures regarding climate risks and opportunities. While public companies have been reporting environmental impacts for some time, they have used different formats with differing information. This has made it difficult for investors and analysts to compare between companies and industries. The development of a standard set of reporting requirements may assist stakeholders in evaluating the environmental impacts on a particular company.

¹ IDC, Seagate Analysts Day 2021, https://s24.q4cdn.com/101481333/files/doc_downloads/2021/2/2021

⁻Seagate-Analyst-Day.pdf

² United States Security and Exchange Commission, "A Climate for Change: Meeting Investor Demand for Climate and ESG Information at the SEC," 3/15/2021, https://www.sec.gov/news/speech/lee-climatechange

Americans support more decisive government action on climate change. A 2020 survey by Pew Research Center³ found two-thirds of adults felt the federal government was doing too little to address climate change. In addition, 52% felt that protecting the environment should be a top government priority. The business world has responded to the challenge of climate change. According to the U.S. Chamber of Commerce, over 90% of the S&P 500 companies now publish sustainability reports. These reports often communicate companies' carbon emissions policies and include environmental impacts such as CO2 emissions, water usage, hazardous materials, and end-of-life product disposal. To support these efforts, the International Standards Organization (ISO) has created a set of standards, the ISO 14001 series, to assist companies in developing, implementing, and managing their sustainability efforts. One of the management tools used is the product Life Cycle Assessments (LCA). This methodology aims to identify all the impacts of a product or service on the environment. Worldwide over 300,000 ISO 14001 certificates have been issued⁴.

Leading IT firms have initiated programs to reduce carbon emissions and improve sustainability. For example, Microsoft has pledged to be carbon neutral by 2030. Amazon announced a goal of reducing carbon emissions by 50% by 2040 and using 100% renewable energy by 2025, five years ahead of the original target. IBM announced a plan to be net-zero greenhouse gas emissions by 2030, including a 65% reduction against their 2010 baseline by 2025. Google has committed to operating on 24/7 carbon-free energy in all their data centers and campuses worldwide by 2030. Other industries such as retail, airlines, and telecommunications have initiated similar initiatives. For example, Walmart's Project Gigaton aims to avoid one billion metric tons (a gigaton) of greenhouse gases by 2030. Southwest Airlines announced its support of the US Department of Energy's National Renewable Energy Laboratory (NREL) effort to develop commercially viable and scalable sustainable aviation fuel. AT&T has committed to be carbon neutral across its entire global operations by 2035.



- ³ Pew Research Center, "How Americans see climate change and the environment in 7 charts," April 21, 2020, https://www.pewresearch.org/ fact-tank/2020/04/21/how-americans-see-climate-change-and-theenvironment-in-7-charts/
- ⁴ ISO, "ISO ISO standards for life cycle assessment to promote sustainable development," July 7, 2006, https://www.iso.org/iso-14001environmental-management.html.

The digital data explosion

Along with society's increasing focus on climate change, another significant development is taking place, the dramatic growth of digital information. This growth is driven by video, medical imaging, video surveillance, autonomous vehicles, the internet of things (IOT), artificial intelligence, analytics, scientific research, expanding internet usage, and 5G networks. IDC estimates⁵ this "Data Sphere" will grow from 33 zettabytes (ZB) of information in 2018 to 175 ZB in 2025. However, only a portion of the information created is stored. They further estimate that approximately 10% of the Datasphere, or about 17 ZB, will be stored in 2025. As large as these estimates are, other advances could push volumes even higher. For example, the Automotive Edge Computing Consortium estimates that globally connected autonomous vehicles could send up to 10 exabytes of data per month by 2025⁶. In addition, IOT could generate as much as 73.1 ZB annually by 2025⁷.

One zettabyte is a massive amount of information. To put this in perspective, a single ZB of storage requires over 55 million LTO 9 18 TB tape cartridges or 55 million 18 TB HDDs. The length of the tape in an LTO 9 cartridge is 960 Meters. If all the tape media in a ZB of LTO tapes was extracted and placed end to end, it could reach Mars (the minimum distance from the Earth to Mars is 54.6 million kilometers).

All the stored data must reside on media of some type. IDC estimates 8.3 ZB of data will be stored in 2021, and 62% will reside on HDDs, 9% on SSDs, 15% on tape storage with the balance on NVM and optical. Different storage types have different technological and economic life spans. For example, HDDs typically offer a five-year warranty and, in most environments, will be replaced at that time. Tape storage is a different matter. Modern tape media has more than a 30-year life. However, while the tape media may last that long, it is usually replaced no later than ten years driven by the economic, technological, and operational benefits of newer media and drives.



- ⁵ IDC, Seagate Analysts Day 2021, page 7
- ⁶ AECC, "Breaking the down the barriers to automotive edge adoption," White Paper, June 2, 2021.
- ⁷ DataProt, "Internet of Things statistics for 2021 Taking things apart," March 24, 2021, https://dataprot.net/statistics/iot-statistics

Storage industry sustainability initiatives

The storage industry has embraced sustainability as a significant challenge, and all the leading providers have established aggressive sustainability goals. For example, Seagate Technology, one of the largest HDD manufacturers, has committed to reducing scope 1, scope 2, and scope 3 Green House Gas (GHG) emissions by 20% by 2025 and 60% by 2040⁸. Western Digital, another large HDD supplier, in their 2020 Sustainability Report⁹, highlights a 44% reduction in the amount of energy consumed per PB in their HDDs. Fujifilm, the leading supplier of LTO tape media, has implemented the Sustainable Value Plan 2030¹⁰. It includes targets for reducing CO2 emissions by 45% compared to the 2013 level and transitioning 50% of their purchased power to renewable energy by 2030.

To address the growing need for transparency regarding sustainability, storage industry suppliers have provided estimates of the CO2 emissions for their products across their product lifecycle. This updated report includes both higher capacity HDD and LTO tape that shipped since the initial 2021 report. For example, Seagate provides a lifecycle CO2e estimate for the high-capacity enterprise Exos X18 hard disk drive¹¹. They estimate that 2.55 kg CO2e/TB/year is generated during its lifespan of 5 years. The LTO tape estimates are based on lifecycle estimates from Fujifilm of the new LTO 9 for CO2e, including CO2 generated by the tape drive energy usage, of .07 kg CO2e/TB/year of CO2e over a ten-year life, 97% less than an HDD¹².



The Life Cycle Assessment (LCA) methodology is often used to develop a complete picture of sustainability. LCA identifies and measures a product's impact on the environment from the extraction of the raw materials used in its manufacture, its use by consumers, and its recycling or disposal. For storage media, this includes CO2 emissions from its usage in the data center but also the associated procurement, manufacturing, distribution, use, and disposal, as shown in Figure 1. Figure 1 – Storage Media Lifecycle

- ⁸ Seagate, FY2020 Global Citizenship Annual Report, https://www.seagate.com/files/www-content/global-citizenship/ en-us/docs/fy-2020-gc-annual-report.pdf
- ⁹ Western Digital, Sustainability Report 2020, https://documents. westerndigital.com/content/dam/doc-library/en_us/assets/public/westerndigital/collateral/company/western-digital-2020-sustainability-report.pdf
- ¹⁰ Fujifilm, Sustainability Report 2020, https://www.fujifilm.com/files-holdings/ en/sustainability/report/2020/sustainability_activity_report_2020_ff_sr_ 2020_all_a4_E.pdf


https://holdings.fujifilm.com/en/sustainability/plan/svp2030/environment

- ¹¹ Seagate, Sustainability Report, https://www.seagate.com/globalcitizenship/product-sustainability/exos-x16-sustainability-report/
- ¹² A discussion of the difference between CO2 and CO2e is in the Appendix.

The benefit of moving cold data to modern tape storage

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. Industry analysts estimate that as much as 60%-80% of all stored information is cold data, yet it often remains on hard disk drives.

Storing cold data on tape storage represents a real opportunity for enterprises worldwide to contribute significantly to society by reducing their carbon emissions and their impact on the environment. To illustrates this opportunity, we calculate the environmental impact in a hypothetical data center that must store 100 PB of data with a ten-year retention requirement. First, we estimate the CO2e emissions resulting from keeping all



100 PB ten years of CO2e emissions

Figure 2 – Ten-Year CO2e emissions

the data on HDDs. We then calculate CO2e emissions of an active archive solution where 60% of the data is placed on tape and 40% on HDDs. We also calculate a deep archive solution where 100% of the information is placed on tape. The analysis includes the CO2e generated due to energy consumption during the media's use and the emissions associated with the acquisition of raw materials, manufacturing, and the final disposal of the storage media. A note of caution, this analysis includes only the emissions associated with the storage media and not the supporting IT infrastructure such as controllers, libraries, tape drives (except energy usage), servers, networks, and supporting infrastructure.

In this hypothetical example, storing 100% of the data on hard disk drives generates 2,663 tons of CO2e 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 1,113 tons of CO2e, a 58% reduction. A deep archive solution storing all the data on tape generates only 79 tons of CO2e, a 97% reduction. Figure 2 illustrates the savings.

In addition to lower carbon emissions, moving data to tape substantially reduces the amount of electronic waste (eWaste). A ten-year retention requirement will require at least one HDD refresh, resulting in

purchasing replacement HDDs in year six. Based on HDD historical capacity growth, the year six refresh is estimated to utilize 38 TB HDDs. The larger capacity HDD substantially reduces the number of HDDs required and the resulting CO2e emissions and electronic waste. Storing all 100 PB on hard disk drives and refreshing them after five years generates 8.2 tons of eWaste. Tape storage has a longer life and does not require a refresh over ten years. As a result, storing

60% of the data on tape results in only 4.0 tons of electronic waste, a 51% reduction. Keeping all the data on tape produces only 1.2 tons of eWaste, an 85% reduction. Figure 3 illustrates the waste reduction.

While the sustainability benefits of storing cold data on tape are substantial, there is also a significant financial benefit. To estimate the economic benefit, using the Fujifilm TCO tool¹³, the TCO for a disk storage solution using high-capacity HDDs is compared to a tape solution using LTO 9 tape media and drives. The estimated tenyear TCO for keeping all 100 PB on disk is \$17,707,468. With the tiered solution, the ten-year TCO drops to \$9,476,339, a 46% savings. Storing all 100 PB of data on





Figure 3 – Tons of eWaste



tape reduces the TCO even further, decreasing the cost to \$3,832,956, a 78% reduction. Figure 4 illustrates the savings. Figure 4 – Ten-Year TCO

¹³ https://www.fujifilm.com/us/en/business/data-storage/resources/tco-tool

Potential Global Impact

Potential sustainability and financial benefits are dramatic for a single firm. But what if organizations across the globe assessed their operations and moved 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 it resides on. Fortunately, IDC, a leading research firm, has developed this estimate in their IDC Global Data-Sphere report.

IDC estimates 8.3 ZB of data will be stored worldwide on all media types in 2021, with 5.1 ZB or 62%, of it residing on HDDs. What if a significant percentage of the diskresident 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. For example, migrating 60% of the HDD to tape reduces CO2e emissions by 79 million tons, a 58% savings.




Figure 5 – Ten Years CO2e with % HDD data moved to tape

Summary

Global warming is an enduring issue. As a result, governments are imposing new and increasingly costly reporting requirements and regulations while consumers favor more robust government policies that protect the environment. To address sustainability, enterprises large and small are using new methodologies to understand and manage the impact of their products on the environment across their lifecycle while implementing a wide variety of programs and initiatives to improve sustainability.

There are many different alternative projects that organizations may adopt to address and adapt to these sustainability challenges. They will vary by industry and company. However, the ideal sustainability program will significantly reduce carbon emissions, minimize a product's impact on the environment, reduce expenses, and be easily implemented. For enterprises with a significant amount of stored data, moving cold data from disk to tape storage is one of those programs. To address this opportunity, the first step is to perform a comprehensive assessment of the disk-resident data to identify "cold data." Once the evaluation is completed, the size of the opportunity will be apparent. Then a project can be organized and implemented that will migrate the cold data from disk storage to tape storage. While not trivial, the software tools and methodologies for managing and migrating large quantities of data have become widely available. By identifying cold data and then migrating the information to tape storage, organizations may significantly reduce their carbon footprint and eWaste while simultaneously reducing costs.

It is also clear that if enough firms worldwide pursued this strategy, the environmental impact would be a meaningful global decrease in CO2 emissions and electronic waste. Given the confluence of events and focus on global warming, clearly, now is an excellent time to evaluate IT data storage strategies and move infrequently accessed data to modern tape storage.



Appendix

Notes on this update

This report revisits the sustainability analysis first published in the summer of 2021 and incorporates the new LTO Generation 9 and higher capacity Seagate Exos X18 drive. LTO tape CO2e/TB emissions improved primarily because the LTO 9 tape cartridge CO2e profile is similar to the LTO 8 tape cartridge, producing 12.98 kg per cartridge over its ten year life versus 13.72 kg/cartridge for LTO 8 (when including an allocation of the tape drive power). However, with an 18 TB capacity for LTO 9 versus 12 TB for LTO 8 the CO2e/TB is significantly lower. In addition, the cartridge weight is identical resulting in fewer tape cartridges for a given

capacity requirement and substantially less e-waste.

The HDD information is updated to include the 18TB Exos drive instead of the 16TB drive. While the number of HDDs required for a given capacity is correspondingly smaller, it is not reduced as much as the LTO 8 to LTO 9 tape media change. Also, according to Seagate, the CO2e/TB/year slightly increased from 2.32 kg-CO2e/TB/Year for the 16 TB HDD to 2.55 kg-

CO2e/TB/year for the 18 TB HDD. The net result is LTO tape CO2 emissions improved relative to HDD since the initial paper.

Methodology

CO2e Emissions

Fortunately, the data storage industry provides excellent data on the environmental impact of the manufacture, use, and disposal of its products. The basis for the hard disk drive CO2e is the Seagate Exos X18 enterprise hard disk drive. The information used in this analysis can be found on their website at this link: https://www.seagate.com/global-citizenship/ product-sustainability/exos-x18-sustainability-report/. Given historical HDD capacity growth rates, it is projected that all the initial HDDs will be replaced in year six with a 38 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 LTO 9 tape media lifecycle CO2e emissions. For a comprehensive picture of carbon emissions, the study added the CO2 emissions resulting from the energy used by the IBM LTO 9 tape drive. A ratio of one drive for every 200 cartridges is used. The packaging and distribution CO2e estimate for LTO is based on the weight of the tape cartridge versus an HDD. The ten-year CO2e emissions total for an LTO 9 cartridge, including its drive energy usage, is 12.98 kg of CO2e per tape cartridge. The tape detail is in Figure 6.

	CO2e -kg
Bill of Materials	1.97
Manufacturing Energy	3.81
Packaging	0.07
Distribution	0.55
Use Phase	5.37
End of Life	1.21
Total CO2 Over Useful Life	12.98

Figure 6 – Estimated LTO 9 CO2e emissions

Tape media suppliers publicize that tape storage has a storage life of more than 30 years when stored properly. However, most organizations will refresh the tape media and drives long before then for practical economic, technical, and operational reasons. This analysis uses a ten-year life.

CO2 and CO2e

CO2 emissions and CO2e are different. CO2 refers to the amount of carbon dioxide emitted into the atmosphere. However, CO2e stands for carbon dioxide equivalents and measures the impact of greenhouse gases on global warming expressed in CO2 equivalents. For example, methane is estimated to have 25 times the effect of CO2 on global warming. Therefore, CO2e is the more comprehensive measure and is used in this paper.

Global Reduction Estimates

FR EE

The potential global reduction in CO2e is based on the quantity of CO2e generated when storing 1 PB of data for ten years on HDD versus tape. For HDD, it is estimated that keeping 1PB for ten years will generate 26.6 short tons of CO2e. For tape, storing 1 PB for ten years will generate .79 tons of CO2e. These rates are then applied to the worldwide quantity of data stored on HDDs in 2021, with various percentages moving to tape to generate the graphic.

Electronic Waste

The estimate of electronic waste is based on the weight of the media. These weights are provided in the technical specifications. For HDDs, a weight of 671 grams per HDD is used; for LTO 9 tape, a weight of 200 grams per cartridge is used. For a given amount of stored data, 100 PB in our example, the required number of units is calculated based on the capacity of the media. The number of units is uplifted for HDDs based on using a data protection method, either RAID or erasure coding, reducing the effective HDD capacity. An effective utilization for HDDs of 75% is used for this analysis.

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 company, product, or service names may be trademarks or service marks of others.

About the Sponsors

FUJIFILM

Value from Innovation

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



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 storage industry economic analysis and

consulting. He started his information technology career with the Data Processing Division of IBM in 1978 and held several 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's in Business Administration and a bachelor's degree in Economics from the University of Arizona. His website is bradjohnsconsulting.com.