

White Paper

Accelerating Green Datacenter Progress with Sustainable Storage Strategies

Sponsored by: Fujifilm

Jennifer Cooke
September 2021

Phil Goodwin

Ashish Nadkarni

Eric Sheppard

IDC OPINION

Improving sustainability of operations is a strategic imperative for business leaders across all industries. IT organizations are in a pivotal position to drive meaningful progress by leveraging technology and implementing processes that decrease waste and optimize the use of all resources. As the pressure builds and the need for proof of progress escalates, progressive and committed organizations will evaluate their storage and compute resources through a lens of sustainability.

Large datacenter operators have been seeking to improve the sustainability of their facilities and infrastructure for more than a decade. As major consumers of energy, datacenters are at the center of conversations around efficiency and CO₂ emissions. Significant progress has been made with the adoption of energy-efficient servers and virtualization technologies to improve utilization rates. Still, energy consumption continues to escalate. Organizations are faced with the need to support the business in a digital-first era in which data creation, collection, and storage continue to grow at staggering rates. Doing this while also finding ways to be more sustainable is a challenge facing all organizations. Addressing the compute and data needs of the business while continuing to improve sustainability requires taking a new look at processes and technologies. Organizations committed to sustainability are seeking to better understand resource usage and to optimize the use of all resources.

Until now, much of the tech industry's and large datacenter operators' focus has been on the shift to cleaner, renewable sources of electricity. Although this is an important aspect of becoming a more sustainable organization, it does not address the most important trait of a green datacenter: the reduction of waste and optimization of resource usage. Renewable energy is not the only path to sustainability, and it's not without drawbacks. A massive opportunity exists to accelerate sustainability progress by reducing the waste of resources and ensuring that maximum value is realized from infrastructure investments. For example, data storage using modern tape systems can accelerate sustainability progress by reducing energy consumption, CO₂ emissions, and electronic waste. A comprehensive and effective sustainability strategy will evaluate all new and existing technologies through a "green datacenter" lens that assesses the total impact on resource optimization.

Pursuing ways to use electricity more efficiently is particularly important for large-scale datacenter and cloud providers. Considering the scale of these facilities, seemingly small changes can make a large difference in the cost for environmental factors and carbon footprint. Much of the focus for green datacenters has been on renewable energy; although cleaner energy sources are an important area of

focus, reducing waste and optimizing resources can also play a significant role in the evolution of the green datacenter. Combining better energy efficiency with cleaner sources of energy offers the best possible outcome for both reduced costs and reduced carbon emissions.

IDC research shows enterprise data storage growing at a 27% compound annual growth rate (CAGR), with data volumes doubling approximately every three years. Nevertheless, not all of this data requires immediate or low-latency access. Much of the data will be infrequently accessed or may lie dormant for years, making it eligible for tiering to much lower-cost, higher-latency media while remaining available when needed. Organizations may choose to retain this data on nearline media or in an “active archive” tier for regulatory and governance reasons, analytics, and other use cases. This data does not need to be kept on hard disk drives (HDD), which require constant power and cooling to maintain operations. IDC believes that progressive organizations that are determined to improve their sustainability will rethink existing data storage practices and consider including more advanced magnetic tape as a practical route to sustainable operations.

Today’s highly advanced magnetic tape has seen a resurgence in the marketplace in recent years, for the following reasons:

- Tape remains, by far, the lowest-cost storage media on a cost-per-GB basis.
- Tape storage itself requires minimal power to operate and does not increase ambient temperatures, thus requiring minimal power to cool the datacenter environment.
- Tape can reliably store data for long periods with an excellent bit error rate and a long archival life exceeding 30 years.
- Tape has emerged as a valuable tool in the fight against ransomware.

Organizations that have moved away from tape or never used tape should consider it for its ability to reduce costs, improve carbon footprint, and provide the greatest possible assurance that data can be recovered from any fault or attack. Even cloud-native solutions are discovering that high-capacity, automated tape systems can play an essential role in the core and cloud.

METHODOLOGY

This initiative measured the impact of tape storage on energy consumption and carbon emissions. For this research, IDC used a conservative and defensible approach to calculate the power consumption that could be avoided by relying more on tape storage. The basis of this analysis is IDC’s Worldwide Quarterly Trackers, which cover technologies including enterprise storage systems. IDC applied assumptions including power consumption by system type and CO₂ emissions per MWh of power to determine the amount of power consumed to operate and cool installed enterprise storage capacity. With this energy consumption data, IDC calculated the amount of CO₂ emitted each year in running and cooling enterprise storage capacity and then calculated the cumulative CO₂ emissions throughout the 2019-2030 forecast period.

IDC then determined the impact of moving more installed enterprise storage capacity to tape libraries to calculate CO₂ savings. This resulted in power savings per year and cumulative power savings as well as CO₂ emissions avoided per year and cumulatively throughout the forecast period to understand the impact of shifting more data to tape storage.

SITUATION OVERVIEW

Large cloud and service provider datacenters are central to everything from business operations to human communication and connection. These facilities are essential resources as the world shifts to digital-first operations, and they are also at the center of conversations about sustainability. According to the environmental, social, and governance (ESG) reporting from several major datacenter owners and operators, energy consumption grew 31%, on average, from 2017 to 2020. Hyperscale datacenter operators are increasing energy consumption at an even higher rate. As these facilities have expanded to multiple exabyte-scale operations, the electricity needed to support a secure and resilient environment is staggering. IDC forecasts that there will be almost 20,000 high-end and mega datacenters around the globe in 2025. If power consumption continues at the current rate, the amount of electricity needed to support these operations will necessitate a new strategy to optimize resource usage.

The irony in the green datacenter conversation is that one of the core value propositions of IT is to improve operational efficiency. By rethinking technologies and processes to align greener datacenter principles with digital-first operations, IDC believes that these power-hungry resources have the potential to break the longstanding formula that has dictated that more data and compute power equals more power consumption.


At the same time that digital-first operations are emerging as highly efficient and using resources wisely, the datacenters that support them are under increasing scrutiny due to the amount of energy they consume. These facilities have expanded to multiple exabyte-scale operations. As workloads and data increasingly become concentrated in large cloud service provider datacenter facilities, the burden of sustainable operations shifts. These providers have made meaningful investments in renewable energy sources and have committed to becoming carbon-neutral in the coming decades. While these investments in greener and cleaner energy are very important, more can be done to improve sustainability and reduce the carbon impact of datacenters. Large cloud providers are leaning on technology to optimize energy consumption.

Data Growth and Impact on Energy Consumption

Shifting to data-driven decisions to improve the customer experience, business resilience, and operational efficiency means that massive amounts of data are being created, stored, analyzed, and moved. The amount of data created is growing at a CAGR of 20-43%. Many large organizations are seeing their data double every two years. Worldwide, the amount of data stored is growing at a CAGR of 27% from 2020-2025 (see **Figure 1**).

FIGURE 1

Global Data Stored, 2015-2025



Source: IDC Global StorageSphere, 2021

Long-term data storage and protection requires physical datacenter resources. These datacenters are designed for resilience and security and require a temperature and humidity-controlled environment. A staggering amount of electricity is needed to power IT equipment and keep datacenter environments cool. Energy consumption varies greatly by location and workload, making accurate energy consumption projections complex.

Understanding total energy consumption from datacenters is complex for several reasons:

- **Lack of reported data on actual consumption:** Many of the large datacenter operators do not report energy consumption figures. IDC observes that some of the leading datacenter operators are embracing greater transparency as a part of their commitment to sustainability. Equinix, Digital Realty, CyrusOne, Facebook, Microsoft, and QTS are a few of the major datacenter operators that are reporting energy consumption.
- **Wide variability of energy consumption by workload type and geographic location:** Bitcoin mining operations have a voracious appetite for energy. Many other workloads consume very little energy. Similarly, a datacenter in a hot, tropical location such as Singapore will require far more HVAC power than one in a cooler region such as Norway. Therefore, it is difficult to make assumptions about energy use by counting the number of servers deployed around the world.
- **Lack of enterprise datacenter transparency:** IDC's survey research uncovered that 25% of organizations do not know how much electricity their datacenters consume.

What IDC does know, from service provider datacenter operators that reported total energy consumption, is that the average energy consumption for multitenant datacenters grew 31% from 2017 to 2020 and an estimated 117% for hyperscale operators.

Future energy consumption will be impacted by server utilization rates and workloads, location of infrastructure in warmer or cooler climates, and the growth of digital operations. The opportunity for

wasted energy is immense. Knowing this, the goal for all datacenter operators should be to maximize the value of energy resources. By focusing on this task and leveraging smarter technologies, great opportunity exists to disrupt the longstanding escalation of energy use and waste.

Renewable Energy Is Just the Start

Energy consumption by the major datacenter owners and operators has skyrocketed in the past four years. Because of their commitment to sustainable operations, these datacenter owners also have invested heavily in renewable energy sources, offsetting their carbon impact by developing or sourcing solar, wind, and hydro power resources. Technology companies are leading the charge with renewable energy investments. In the United States alone, eight of the top 10 investors in renewable energy are technology companies, showing the commitment that the sector has to driving meaningful change. This momentum is building out renewable energy sources for municipalities and supporting an ecosystem for broader adoption, making cleaner energy accessible to many others.

Whether through direct usage or, more frequently, renewable energy offsets or credits, the shift from coal and natural gas-based to solar, wind, and hydro power has been the focal point of greener datacenters. By leveraging renewable power when practical, large cloud and hyperscale datacenter operators have reduced their CO₂ emissions and have begun their journey to 100% carbon-free in the coming decades.

But it's not enough to invest only in renewable energy, for several reasons:

- Renewable energy can't come online fast enough or cheaply enough to keep up with the rapid growth in datacenter power consumption to support digital transformation.
- Resilience is top priority as new processes and operations shift to digital-first. Having electricity that is not dependent on sun, wind, or water is a necessity as datacenters support our digital world.
- Solar and wind-power infrastructure have their own negative impact on the earth's resources: Building, recycling, and disposing of solar panels and wind turbines has an associated impact on the environment as well.

Today, a relatively small percentage of the overall energy used by datacenters can be supplied from renewable resources. Therefore, energy conservation and the optimization of energy use must also be the focus of every large datacenter operator. By placing an equal focus on reducing energy consumption and wasted energy, large datacenter operators can demonstrate measurable progress on the sustainability journey. Approaching every technology through this lens, tape storage emerges as a viable, cost-effective way to reduce energy consumption. At the same time, it can improve resiliency through data security.

Accelerating the Sustainability Journey with Green Datacenter Technologies

Maximizing resource usage is one of the key value propositions of infusing IT resources and capabilities into business processes. In the same way IT can make a business run "smarter," IT can make datacenters run "greener." Using technologies that monitor environments in real time, improve IT resource utilization, and automate workloads, organizations can dramatically reduce energy consumption. The following are some examples of how "green datacenters" leverage technology and new processes to change the equation that more compute power and storage equals more energy consumption:

- Investing in AI-influenced cooling control to make many small adjustments that result in significant energy savings (Google reports that it is saving 5% on its energy bill in this way.)
- Rethinking data storage technologies and developing a strategy to understand and choose the most sustainable method for data retention, protection, and retrieval
- Creating highly agile environments that enable workload shifting to take advantage of renewable energy sources in real time
- Improving IT utilization rates through virtualization and containerization technologies (By improving server utilization rates, organizations are not only deploying less equipment but also keeping equipment installed for a longer period of time, reducing ewaste and maximizing asset life cycles.)
- Increasing transparency into resources with real-time monitoring and control

These strategies have the potential to accelerate sustainability progress by reducing the energy required to power compute and storage resources and cool the environment in which they reside.

Why Tape Should Be Included in Every Green Datacenter Strategy

Effective sustainability strategies should be measured by their ability to reduce waste, optimize use of resources, and contribute to reducing greenhouse gas emissions. Organizations that are developing greener strategies may believe that every sustainable technology needs to be “new” when, in reality, an effective strategy is to evaluate each technology for the impact it has on CO₂ emissions. Savvy CIOs and leadership are tuned in to “circular economy” principles and understand that just because a product is new does not mean it is “green.”

Tape supports green datacenter initiatives by storing data offline without consuming energy. Tape systems can also serve as nearline or active-archive tiers of storage. This means that the data is available for productive use but does not consume processing power and requires minimal environmental resources while the data is not being used. Organizations have learned that larger data sets lead to better analytics, and active-archive tiers can increase the data set size while making the data available without human intervention.

Capacity to Support Future Growth, Sustainably

The recently released Generation 9 of LTO Ultrium tape features a native capacity of 18TB (45TB compressed). The linear tape open (LTO) product road map currently extends to Gen 12, which anticipates 144TB (360TB compressed) per cartridge. The road map for LTO tape continues to significantly increase areal density (i.e., bits per square inch), meaning that a single tape cartridge will store greater amounts of data with each new generation, thus improving its low TCO and CO₂ profile.

Other Benefits of Tape

Modern magnetic tape has also become an invaluable tool in thwarting ransomware, data theft (exfiltration), and other malware. Specific capabilities of tape include:

- **Immutability:** Tape can keep immutable copies of data, sometimes called WORM (write once, read many), that cannot be altered or deleted. Because ransomware hackers often attack backup copies to encrypt, scramble, or delete them, immutable copies are critical to assuring data integrity and recoverability.
- **Encryption:** Encrypted tapes ensure that people not authorized to read the data on the tape cannot do so. Thus, tape theft will not benefit criminals wishing to steal the data, whether it is an internal or external attempt to do so. Encryption of tape data is a best practice and is

especially critical for tapes that are moved from one location to another. Because tape systems have the ability to encrypt data at the hardware level, encryption can be implemented without application or processing overhead.

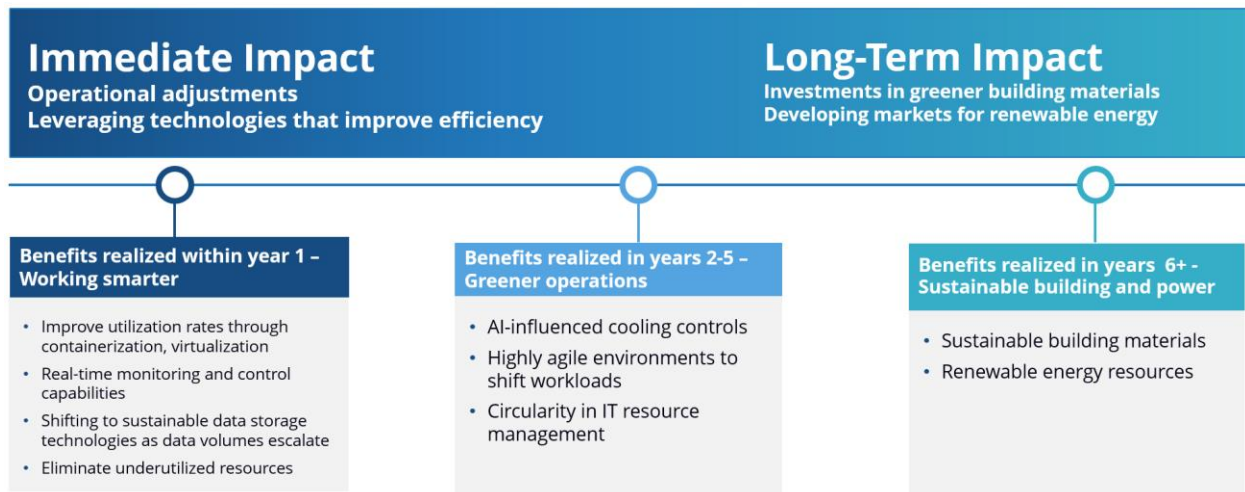
- **Air gap:** As noted above, hackers often attack backup copies to make recovery impossible. A reliable means to defeating such an attack is by creating a physical disconnect (“air gap”) between the primary and backup copies, to prevent the criminals from physically accessing the backup. For tape, this is as simple as removing the tape from the tape library. Some organizations may choose to move the tape offsite for additional protection. The best practice is to retain one offline copy onsite and another offsite. While it is possible to create air gaps with online systems, tape is far and away the easiest, most reliable way to implement a true physical air gap.

Green Datacenter Outcomes

Green datacenter technologies can support both longer-term and immediate impacts on CO₂ emissions, resource optimization, and cost reduction (see **Figure 2**). Shifting to renewable energy is a longer-term investment and builds out infrastructure that supports a pivot to totally new energy generation processes. These investments will have broader impacts on communities and the grid as datacenter investments help build an ecosystem and infrastructure for widespread deployment. Major datacenter owners and operators, as major consumers of electricity, are leading the way in renewable energy infrastructure buildouts. Shorter-term outcomes are realized by deploying technologies such as AI-assisted cooling, energy efficient data storage (tape), virtualization/containerization, and real-time monitoring and control capabilities.

FIGURE 2

Near- and Long-Term Benefits of Green Datacenter Technologies



Source: IDC, 2021

Leveraging Tape Storage to Accelerate Green Datacenter Initiatives


When making decisions on data storage infrastructure, consideration should be given to the long-term impact on energy consumption. To better understand the impact that tape storage could have on CO₂ emissions, IDC performed an in-depth analysis and created a scenario to understand the impact if more data is moved to tape storage. When taking into account the total resource usage over the life span of data storage, shifting to tape can result in a significant, measurable difference in energy resource consumption. The energy cost savings as well as the CO₂ emissions reduction are compelling reasons to consider greater use of tape storage.

Expanding the use of tape storage can positively impact sustainability goals. The positive impact can be recognized immediately, with reduced electricity consumption. In a scenario in which an increasing amount of data to be stored is designated as “archival” and 80% of archive data to be stored on enterprise storage systems and 57% of replicated data to be stored on enterprise storage systems is migrated to tape, the annual CO₂ reduction by 2030 is 43.7% (see **Figure 3**). The opportunity to positively impact the environment by shifting to tape is staggering, with the potential to avoid a cumulative 664 million metric tons of carbon emissions between 2019 and 2030. This amount is equal to the greenhouse gas emissions from 144 million passenger cars driven in one year or the amount of energy used by 80 million homes in one year.

The opportunity to positively impact the environment by shifting to tape is staggering, with the potential to avoid 664 million metric tons of carbon emissions.

FIGURE 3

Worldwide Power Consumption/CO₂ Emission Avoidance Associated with Shifting to Tape Storage



Source: IDC, 2021

FUTURE OUTLOOK

As organizations become more sophisticated in their green datacenter and sustainability strategies, they will look for ways not only to shift to more renewable sources of energy but also to use less energy overall.

In Europe today, executive leadership is compensated based on their ability to meet sustainability goals. IDC believes that in the coming years, this practice will be implemented in North America and other regions as well. Such incentives will increase interest in new ways to understand and create greener datacenter resources. Investors, top talent, and ecosystem partners will increasingly depend on the ability to prove that progress is being made.

As these proof points become competitive differentiators, organizations will also come under much more scrutiny to provide data not only on the percentage of clean energy being used but also on the total amount of resources, including electricity, being used. The momentum for a greener, more sustainable world will pressure organizations across many areas of their business and deep into their supply chains. For this reason, IDC believes it is a strategic imperative for organizations to expand their sustainability initiatives to focus on the improved use of all resources. Savvy organizations will reexamine all processes based on the ability to protect the earth's environment. As this practice becomes more commonplace, another benefit that organizations will realize is the potential for cost savings and greater profitability by implementing solutions such as tape storage that reduce both costs and energy consumption.

About Fujifilm

Fujifilm is a worldwide supplier of magnetic tape cartridges and other technology solutions. The company provides datacenter customers and enterprise industry partners with a wide range of innovative recording media products and archival solutions. Based on a history of thin-film engineering and magnetic particle science such as Fujifilm's NANOCUBIC and Barium Ferrite technology, Fujifilm is an innovator in data storage products. Worldwide, Fujifilm and its affiliates have surpassed the 170 million milestone for the number of LTO Ultrium data cartridges manufactured and sold since introduction, establishing the company as a major player in the midrange and enterprise tape market.

The Fujifilm group is committed to responsible environmental stewardship and good corporate citizenship. It issued its environmental, social, and governance (ESG) plan "Sustainable Value Plan 2030 (SVP 2030)" in 2017. The plan targets FY 2030 as a long-term goal and is in accordance with the Paris climate agreement. This plan includes six priorities: environment, health, daily life, work style, supply chain, and governance. The company emphasizes the environment as most important, setting a target for reducing CO₂ emissions for the entire product life cycle, from procurement to manufacturing to distribution and disposal, by 45% compared to 2013.

Challenges/Opportunities

- Tape is often overlooked as a modern technology, yet it has kept pace with the growth of data while delivering better cost-per-GB ROI with each generation. Many buyers overlook tape without understanding or considering the value that modern tape systems deliver to the datacenter. All tape vendors, including systems and media suppliers, are challenged to educate buyers appropriately and change misconceptions.
- Simply adding tape to the datacenter is not the solution, as IT organizations need to architect it into their overall data management strategies.

- Properly done, tape is another tier of storage with its own service level standards for data availability. Organizations may reflexively retain data on HDD without considering either the media cost or the carbon emissions implications. Organizations that implement tape nearline and in active-archive tiers can have all of the benefits of data retention while minimizing the per-GB storage cost and contributing to a significant reduction in carbon emissions.
- As a tape media supplier, Fujifilm is only part of the solution. The company must work with its supply chain and industry partners to foster a complete solution.

CONCLUSION

Understanding technologies that can support green datacenter operations is a critical strategy for CIOs and executive leadership. Investors, ecosystem partners, and future talent will be increasingly sophisticated in their assessment of what constitutes a truly “green” datacenter. They will seek data and metrics that offer proof of progress and demand transparency in reporting. Shifting to tape storage can accelerate sustainability progress by significantly reducing energy consumption. The direct reduction in consumption has an immediate and positive impact on CO₂ emissions by reducing wasted energy resources and overall consumption. By focusing on ways not only to shift to greener and renewable energy sources but also to reduce overall energy consumption, organizations can show true progress on sustainability initiatives.

Company leaders should evaluate all IT resources through a lens of sustainability – one that ensures minimal waste and supports greener operating principles. Doing so will uncover the measurable benefits of shifting more data storage to tape. Based on IDC’s analysis and comparison of energy consumption trends, by leveraging more tape storage, there is an opportunity to reduce CO₂ emissions by 43.7% and 664 million metric tons. Large datacenter operators understand the critical need to be more sustainable and environmentally responsible. IT leaders have the opportunity to disrupt the equation that more compute and data equals more energy consumption. Progressive, sustainability-minded IT leaders can dramatically reduce carbon emissions by rethinking their data storage options and leveraging tape storage where appropriate. The key for these organizations is developing a framework and process to demonstrate that their strategies are contributing to positive change. Tracking energy consumption, data growth, and performance metrics is essential to understanding the impact of changes.

Organizations that are able to deliver transparency and metrics on sustainability progress will be more competitive. Investments made today in technologies that reduce energy consumption and optimize resources will not only reduce spend and waste but also improve opportunities to attract investors, top talent, and strategic ecosystem partners. More importantly, shifting to tape and reducing electricity consumption has the potential to dramatically reduce the amount of CO₂ emitted, thus positively impacting the earth and supporting important sustainability initiatives.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2021 IDC. Reproduction without written permission is completely forbidden.

