White Paper

Implementing Sustainable Storage Infrastructure

Sponsored by: IBM

Ashish Nadkarni Bjoern Stengel
January 2023

IDC OPINION

Sustainability is an important topic in the C-suite agenda. Enterprises are consistently increasing their IT infrastructure footprint to drive product and services innovation, accelerate routes to market, and modernize business operations. But something else is happening in tandem. Saddled with burgeoning energy costs, CXOs are looking at ways to make their IT infrastructure investments more power efficient. The governmental and social mandates to be socially and environmentally responsible put additional sustainability dimensions on these investments.

Historically, within the IT infrastructure domain, the focus of data storage investments has been on security, performance, scalability, availability, and manageability. Unlike compute requirements, which can vary from workload to workloads, storage systems are foundational and serve the needs of enterprisewide data access. With 84% of enterprises actively pursuing digital transformation, storage requirements have evolved. Data growth rates average at least 30–40% per year for many enterprises, which today find themselves often managing multi-petabyte data sets. This storage infrastructure must be agile, more efficient, and deliver consistent performance at scale. The increasing importance of IT infrastructure to operational decisions has also raised availability requirements, which present their own challenges at multi-petabyte scale. And powering datacenters operating at this scale requires better advance planning than ever before.

A 2021 IDC datacenter operations survey revealed that one in four United States–based IT organizations had reported delays in IT deployments due to power or space constraints. In the past, power efficiency has not been high on most IT decision makers' (ITDMs') list of purchase criteria primarily because power costs were paid for out of facilities management rather than IT budgets. When IT projects are delayed, however, because of power constraints, it becomes an IT issue. For ITDMs whose datacenters are at risk of running out of available power as their IT infrastructure grows, the issue is not whether a system can be built that meets performance and capacity requirements. It is whether such a system can be built within the available power envelope. IDC believes that within two to three years, power efficiency will become one of the top 5 criteria for enterprise IT infrastructure purchase decisions for enterprises with larger on-premises IT infrastructure requirements.

Most enterprises that are serious about storage infrastructure efficiency have already moved to flash for primary (i.e., mission-critical, latency-sensitive, and highly available) workloads. The next phase of power efficiency at scale will come from IT vendors like IBM that are making sustainability at the core of the design principles for delivering highly efficient vertically integrated storage solutions that can meet or exceed the sustainability targets of a digital-first enterprise.

January 2023, IDC #US49980323

SITUATION OVERVIEW

Sustainability Is a C-Suite Agenda

Sustainability has become a mainstream business topic, one that is at or near the top of a C-suite agenda at most enterprises globally, across every industry. Environmental, social, and governance (ESG) approaches have evolved significantly over the past few years, fueled by interest from various stakeholder groups, including investors, regulators, customers and consumers, and employees. There is now a profound sense of urgency of taking steps to ensure the longevity of the planet thrusting the corporate world into the center of action.

The promising news is many organizations are responding to the pressure to act by looking at sustainability primarily as a competitive differentiator and secondarily as a cost or reputational risk factor. This enables them to infuse sustainability into their strategic decision-making process and objectives to increase their operational and financial performance. As sustainability becomes an integral part of an organization's business strategies, CXOs are applying a more deliberate, materiality-driven lens on ESG topics that have the greatest impact on their organization's enterprise value.

Corporate views about sustainability are evolving rapidly, and organizations are incorporating these views into their business practices, including procurement. Now that many organizations have an environmentally focused mandate, we are seeing input from purchasing, line of business (LOB), and IT as they develop preferences and standards for equipment purchases that meet standards for all stakeholders in the public and private sectors.

Sustainability and IT Infrastructure

ESG topics – and environmental sustainability in particular – have several implications for ITDMs and CIOs. IDC's datacenter research paints a sobering picture of how datacenters have become a major source power and water consumption globally (see *Datacenter Sustainability,* IDC #US48153021, August 2021).

Datacenter Power and Water Consumption

- IT represents a significant percentage of U.S. energy use. Datacenter usage of electricity will increase but at a lower rate due to improved efficiency. In 2020, U.S. datacenters consumed approximately 73 billion kWh.
- In 2020, U.S. datacenters used 660 billion liters. This is split across two main categories: electricity generation and cooling. Water use in electricity was x4 greater than that used onsite for cooling: 7.6L of water is used for every 1kWh of electricity generated compared with 1.8 liters per kWh of total datacenter site energy use.

CIOs and ITDMs are therefore directly involved in developing and meeting or exceeding their organization's sustainability goals, especially with datacenter power consumption and efficiency in mind. As their organizations accelerate their digital transformation initiatives, the efficiency and carbon footprint of IT infrastructure are being put under a magnifying glass – but they are not looked at independently but rather as two interlapping areas. Two key proof points that highlight the increasing influence of this trend, via IDC's 2021 *Sustainability Buyer Value Survey*:

- 83% of the respondents agreed that sustainability plays a "very important" role in the planning and procurement decisions for IT equipment in their organization.

- 51% of respondents worldwide already recognize that their sustainability efforts deliver operational and financial benefits.

Among the other business value improvements these organizations realize are enhanced brand perception and increased customer loyalty. These improvements show that organizations are deriving value that goes beyond their investor-focused environmental, social, and governance reporting requirements.

Datacenter power consumption has become an IT responsibility. Thus any constraints put on the rapid deployment of digital transformation projects immediately raise the question of power efficiency of IT infrastructure equipment. Power efficiency is thus gaining priority on the list of purchase criteria for IT equipment, and specifically storage. It requires IT organizations to reconsider their approach to building out the required infrastructure to extract scale without compromising efficiency.

Sustainability and Storage Infrastructure

Storage is a crucial and dominant component of the IT infrastructure stack. If data is like water, maintaining its potability requires a scalable, performant, and highly available storage layer. Without a persistent storage layer, data is ephemeral at best, no matter how capable the compute layer. Historically, this has been the top purchase criterion for enterprise storage infrastructure — to deploy a storage environment that can provide security, performance, scalability, availability, and manageability in line with the rest of the infrastructure.

Digital transformation has further refined the role played by storage infrastructure. The focus of enterprise storage investments has shifted toward efficiency, agility, and delivering consistent performance at scale for new and emerging workloads like artificial intelligence. The increasing importance of IT infrastructure to operational decisions has also raised availability requirements, which present their own challenges at multi-petabyte scale.

Enterprises often struggle between the choice of hard disk drives (HDDs) and flash when designing their storage infrastructure. On a capacity and raw dollar-per-gigabyte cost of capacity basis, HDDs are still superior then off-the-shelf solid state disks (SSDs), but their differentiation ends there. When it comes to efficiency, the story dramatically changes. Here, flash offers the highest efficiency in storage infrastructure for primary (i.e., mission critical, latency sensitive, and highly available) workloads. In fact, the performance characteristics of hard disk drives impose serious limitations when trying to create storage infrastructure as efficient as that which can be built using flash media for performance-sensitive workloads.

Designing Highly Efficient, Performant High-Capacity Storage Systems

With advances in storage media technology, the density of storage devices is on the increase. 15TB SSDs are commonly available from many enterprise storage vendors, and even 30TB SSDs are available for use with many enterprise storage systems. While one might infer that increasing media density should allow IT managers to increasingly handle higher-capacity requirements more cost effectively, and other advancements with persistent and storage-class memory are opening opportunities to build new, higher-performance storage architectures than what exist today, it is not necessarily that simple.

Designing performant high-capacity systems requires vendors to go beyond just assembling a storage system with off-the-shelf components. These vendors design their own proprietary storage devices that give them an advantage over those rely on off-the-shelf SSDs:

- They manage flash media directly rather than delegating media management in each storage device to a device-resident controller.
- They tend to manage flash globally across all devices, making read, write, media refresh, free space management (often referred to as garbage collection), and other media management decisions based on global visibility.
- They build devices that can support significantly higher capacities than off-the-shelf SSDs.
- They provide more "channels" within each device to access these larger flash capacities.
- Even in the face of well-accepted standards, proprietary devices can provide meaningful, sustainable differentiation.

Over time, as a technology becomes more widely adopted, there is often market momentum toward industry standards. When standards meet customer requirements, they make it easier for a particular technology to get to volume production, which makes it easier for customers to integrate it and generally leads to lower prices.

These architectural design points enable higher-density media access to be much more optimally scheduled, avoid bottlenecks that off-the-shelf SSDs cannot and, in general, make much more effective use of available flash capacity. Highly parallel protocols like NVMe offer the ability to leverage massive parallelism when accessing media, taking advantage of these extra channels. The difference in parallel access that can be supported by an NVMe-based proprietary storage device is significantly greater than that which can be supported by a SCSI-based off-the-shelf SSD.

Many vendors still view enterprise storage as a commodity product, as is evidenced by the relatively low levels of R&D investment they make in it (around 5–7% of sales). For these vendors, it's difficult to go the proprietary route (they can't afford it), but they may have also chosen the commodity route for other reasons (like the perception that they meet customer requirements just fine). Commodity off-the-shelf storage devices deliver good performance and capacity and are used by many of the leading enterprise storage providers in the market (as measured by revenue-based market share).

However, vendors like IBM do not use off-the-shelf SSDs, instead building their own proprietary storage devices. These vendors justify their decision by claiming that by building their own devices, they can deliver better performance, higher storage density, reduced power consumption, longer media endurance, and lower cost (at the system level, not necessarily at the individual device level). These are a major contributor to differentiating power and cost efficiency at the system level.

SUSTAINABILITY AND IBM

Sustainability has been a central focus to IBM's Systems business for a long time now. It has had extensive experience, practical outcomes, and internal circular economy policies, enabling the firm to be one of the leaders in the storage industry in terms of sustainable practices.

IBM released its initial environmental policy in 1971. Since then, it has demonstrated unwavering dedication to providing honest and comprehensive information on its environmental impact. For more than three decades, IBM has been issuing an annual environmental impact report, demonstrating its commitment to environmental openness. IBM currently offers transparent access to its progress

toward its 2030 target of producing zero net greenhouse gas emissions. Using the Product Attributes to Impact Algorithm (PAIA) approach created by the Massachusetts Institute of Technology (MIT) Materials Systems Laboratory and its partners, IBM maintains this history of openness by making its Product Carbon Footprint (PCF) reports for its storage devices available to the public.

As a corporation, IBM achieves sustainable leadership by implementing circular economy sustainable policies throughout the organization. There are three key aspects to IBM's circular economy policies — IBM's Design for Environment (DfE), first-tier supplier management, and product waste policies:

- Since 1991, IBM has mandated that new ideas use environmentally friendly materials, increase energy efficiency, and may be readily disassembled and recycled once their useful life has ended (EOL) in accordance with the principles of "Design for Environment."
- Since 2010, IBM vets all its first-tier suppliers through its supplier management system to ensure they meet IBM's corporate sustainability supply chain standards. For example, an outcome of this on FlashSystem is that 99% of paper- and wood-based packaging used is from sustainably managed forests.
- IBM Global Asset Recovery Services (GARS) achieves the highest standards in the market around sustainably handling product waste. IBM recycled, repurposed, or resold 97.7% of its product waste in 2021. Customers can rest certain that when IBM's goods reach the end of their useful lives, they will be disposed of in a safe and environmentally responsible manner.

Over the past few years, IBM has also been building out its portfolio of sustainability solutions to help enable its customers' sustainable transformation. This includes Envizi, an ESG performance management software platform, as well as the IBM Environmental Intelligence Suite to gain insights regarding climate and weather impact, among other components.

Figure 1 illustrates IBM's product end-of-life-cycle processing methods. A breakdown of IBM's product end-of-life 2021 processing methods shows that only 0.3% of the materials in a machine end up in landfills. This is achieved through rigorous system design and production methods:

- From a supply chain point of view, all first-tier suppliers must adhere to social and environmental responsibilities.
- When it comes to manufacturing, a flexible design supports ongoing energy savings, and nearly 60% of the electricity consumed in 2020 was from renewable sources.
- The machines ship in reusable wooden crates, and the wood-based packaging comes from certified sustainably managed forests.
- Maximize the use of recycled and recyclable materials.

FIGURE 1

Product End-of-Life Processing Methods, 2021

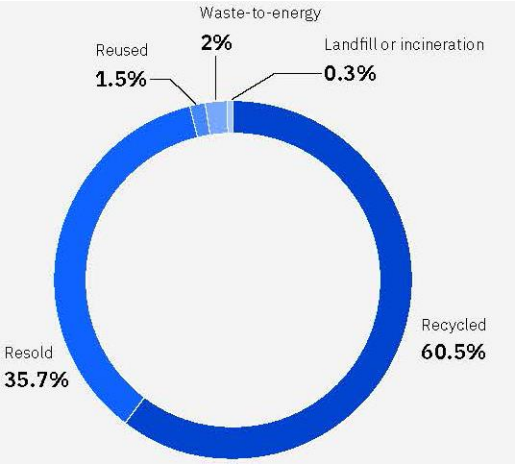

Source: IBM's ESG Report, 2021

IBM FlashSystem Design

IBM has translated its focus and history on sustainability into differentiation for the FlashSystem product line. When considering sustainability in the storage industry, it is important to consider four key factors:

- **Power consumption.** How nimble is the system at consuming power, and do the vendor's choices make a material difference?
- **Energy efficiency.** How dense is the system, and how well does it allocate consumer power to usable capacity?
- **Carbon footprint.** What is the electronic waste profile of the system, what is its carbon footprint, and how much footprint can be "recycled" by making use of existing hardware?
- **Product life cycle/disposal.** It includes the vendor's sustainability goals and transparency in meeting or exceeding them.

Power Consumption and Energy Efficiency

Providing customers with the ever-increasing storage capacity they require while reducing energy use is a formidable challenge. To meet the growing demand, IBM has focused on improving energy efficiency, performance, and access density.

IBM has overcome the resulting complication efforts to improve energy efficiency with innovative engineering efforts. This results in IBM FlashSystem having industry-leading power per rack unit and high-capacity density, which allows clients to have the most energy efficient solutions that require less space, power, and cooling in datacenters. IBM claims that IBM FlashSystem has 122% higher raw capacity per rack unit in the controller compared with the next highest competitor. Key reasons for this density start with the system design.

IBM's FlashCore Module (FCM), a form of computational storage, shifts power-hungry processor (CPU) operations to more efficient solid state drives. IBM has had transparent compression in the FlashCore module since the inception of its end-to-end NVMe FlashSystem product line in late 2018. By eliminating the need to conceptually remap the address space and perform costly garbage collection, this does more than just offload the compression itself; it also simplifies the software storage functionality. IBM has spent significant efforts to improve the FCM garbage collection algorithms for real-world use situations.

Offloading the compression processes to the FCM enables IBM to make use of less resource-intensive storage controllers in the FlashSystem, thus reducing its the energy footprint. A simplified garbage collection process that is performed closest to the NAND flash media also translates into energy savings since metadata and data itself does not have to be moved out of the SSDs into system memory for processing.

IBM FlashSystem noteworthy stats

IBM FlashSystem products are designed to meet IBM's sustainability goals while being very competitive in performance per energy metrics:

- 30% more effective storage capacity per rack unit with FlashCore Module Generation 3
- 809% higher transactional IOPS/watt compared with competitors (on average)
- 122% higher in controller raw capacity per rack unit compared with next highest competitor product (on average)
- 99% of paper- and wood-based packaging used coming from sustainably managed forests

IBM has also invested in making QLC Flash suitable enterprise business settings. The FCM in every model of the FlashSystem family is composed primarily of QLC. There is no option to select separate tiers for capacity and performance. IBM's work enables the use QLC without endurance limitations and with excellent performance.

Carbon Footprint and Product Life Cycle

To assess their overall carbon footprint reduction objectives, it is necessary for vendors to take a system view rather than individual parts.

As an example, on paper, QLC does not inherently use less energy than TLC. In fact, it takes longer to program, which means that the power consumption can in fact increase slightly. While TLC flash tends to be more compact, QLC tends to be much denser. Since February 2020, IBM's FlashCore Module has supported a 38TB capacity. This was more than twice the highest capacity of any other off-the-shelf NVMe drives and is still considered to be the densest in the industry. This density allows IBM to design storage systems with less rack space and reduced cooling requirements, thus reducing the total energy spend. The current third-generation FlashCore Module has 30% more effective storage per rack unit compared with the previous generation. Since critical data services functions are offloaded to the FCM, the storage system offers the needed performance for demanding workloads

without compromising energy efficiency. The resulting product is a compact storage system that requires less cooling while providing exceptional performance in terms of watts per rack unit.

Reducing carbon footprint is only part of being environmentally responsible. System longevity and lifespans are essential to reduce extensive and costly upgrade cycles and ewaste. IBM FlashSystem products also help clients reduce product waste by offering market-leading capacity density, high utilization rates, and longer product lifespans.

IBM's FlashSystem design goals ensure speed and functionality enhancements without the need for powerful (and more energy-hungry) storage servers. IBM delivers it in two ways:

- Software enhancements for existing storage controllers minimize the need for controller swaps. Further, IBM's Spectrum Virtualize software enables IT organizations to upcycle their current devices, utilizing them more and extending their lives.
- Using a combination of predictive technologies with Storage Insights, health binning, and variable voltage technology on their transistor gates, IBM's FlashCore Modules offer remarkable longevity.

Because of these enhancements, the FCMs have up to seven times the flash durability of commonplace commodity SSDs. IBM has also included ample provision in the storage systems to permit the clients to add FCMs to grow even by a single FCM at a time based on requirements.

ESSENTIAL GUIDANCE FOR IT DECISION MAKERS

If ESG considerations are important to the enterprise, ITDMs must lay a good foundation for understanding where power and cost efficiency rank in terms of IT infrastructure purchase criteria by reviewing the relevant regulatory requirements and the company's ESG policies.

ITDMs must view ESG performance as a top 3 decision factor for IT equipment purchases, ensuring that storage-related RFPs include metrics regarding carbon emissions, material use, and labor conditions. In the context of IT infrastructure purchases, RFPs must also seek to determine the vendor's waste management practices and the circularity of value chains as well as the social impact of those value chains (e.g., responsible sourcing practices of raw materials and manufacturing that pay attention to the avoidance of human rights violations, child labor, negative community impacts, and so forth).

IDC believes that in two to three years, power efficiency will become one of the top 5 criteria for enterprise IT infrastructure purchase decisions for enterprises with larger on-premises IT infrastructure requirements. For ITDMs whose datacenters are at risk of running out of available power as their IT infrastructure grows, the question that must be asked is not whether a system can be built that meets performance and capacity requirements but whether such a system can be meet those requirements within the available power envelope.

When evaluating systems based on off-the-shelf SSDs versus those that use proprietary storage devices, IT should pay particular attention to the level of utilization that can be supported while still meeting performance service-level agreements (SLAs). Although the level of utilization of systems depends on several variables, systems that more efficiently use available flash capacity will be able to store more data for a given amount of raw capacity. If, for example, a proprietary storage device-based system can meet performance requirements while being operated at 75% utilization while an off-the-

shelf SSD-based system can only be operated at 65% utilization, that 10% translates into additional storage capacity that does not have to be purchased.

If claims about media utilization rates are true (and any vendor claims in this area should be validated), then, even at just the several hundred terabytes scale, it means that the use of off-the-shelf SSDs will require the purchase of significantly more storage devices to meet a given data storage requirement, a hidden cost that is difficult to identify by just looking at sales quotes for a given amount of raw capacity.

While staying within the stated power envelope must be an overarching consideration, ITDMs must consider the incentive of cost savings associated with purchasing more power- and cost-efficient storage infrastructure. There is a real difference in the power and cost efficiency of different architectures, just as there is a real difference in scalability.

ITDMs should not overlook the impact on total cost of ownership of purchasing more power- and cost-efficient systems. In any bake-offs that may be conducted do as part of a storage purchase, they must make an accurate energy, floorspace, and cost comparison based on actual configuration required for the organization. Larger devices sizes may lower the costs but may not meet performance requirements needed for the environment.

Further, they could lead to higher device counts and thus infrastructure costs. ITDMs must also examine flash capacity utilization differences in different storage system architectures. For example, a 10% difference in utilization means having to buy an extra 100TB of capacity for a petabyte-size storage system.

In conclusion, ITDMs must gain a clear understanding of the performance and capacity objectives for the system being acquired. Building a high-capacity system is very different from building a performant, high-capacity system. If the design goals call for more performance, they should examine systems that can deliver it at scale. Otherwise, they'll end up paying for performance they don't need.

OPPORTUNITIES FOR IBM

IDC believes that as global commitment to sustainability grows, accountability will increase through fines or compensation. Ideally, sustainability officers (SOs) and CSOs will encourage adoption of circular economy principles, which can deliver several benefits: demonstrate environmental commitment and leadership, strengthen SITAD programs, bolster brand reputation, and improve metrics associated with reuse, recycle, and refurbishment (see *IBM: Enabling Secure IT Asset Recycling, Refurbishment, and Resale Through Global Asset Recovery Services,* IDC #US49585122, August 2022).

IBM has nearly four decades of experience in IT recycling, refurbishment, and resale. The company is recognized as a leader in environmental management, voluntary environmental initiatives, product stewardship, sustainability, and ewaste reduction. In the 1990s, IBM spearheaded the development of cross-industry guidelines for voluntary corporate reporting. These guidelines have since become common business practices. In addition, the recent acquisition of Envizi, and the integrations of the GARS team and Technology Lifecycle Services group into the IBM Systems umbrella, indicates a renewed focus on this important area. To further strengthen and differentiate GARS, IBM may want to consider:

- Communicating the consulting services associated with asset life-cycle management and how the reorganization will help its customers achieve its sustainability initiatives
- Clarifying the benefits of the sustainability and enterprise platform, specifically with respect to ease of use and reports that simplify compliance efforts
- Going beyond reporting GARS operational data and describing sustainability stories and outcomes

On the Systems side, the opportunity for IBM lies in gaining mindshare among ITDMs for whom sustainability is an important factor when making infrastructure purchases. There are two areas where IBM must continue its efforts for market awareness:

- With FlashSystem, IBM has made significant advances in delivering an innovative, highly scalable, agile, and efficient storage system. It is part of IBM Systems — one of the key business units for the vendor, one in which it has continued to invest. IBM Systems is backed by IBM Research, which is one of crown jewels of the company. Converting R&D know-how from research to products is something IBM has mastered over several decades.
- IBM's sustainability initiatives run deep within the DNA of the company. It is not superficial to just ensure a checkbox when delivering infrastructure solutions. IBM released its first environmental policy in 1971, decades before sustainability was introduced as an IT topic. It is one of the few leading IT vendors that maintains its history of openness by making its Product Carbon Footprint reports available to the public.

CONCLUSION

Enterprises are encouraged to evaluate power and cost efficiency differences more effectively between different vendor offerings, identifying key metrics that will become important as IT organizations look to modernize to more efficient IT infrastructure. Considering developments just within the past six months, enterprises are changing their views on the importance of environmental, social, and governance issues in making storage infrastructure purchases. Given the complexity of factors that impact power- and cost-efficiency considerations, it can be difficult to understand how and why different storage architectures and product strategies perform on these metrics. The purpose of this document is to help IT managers making storage infrastructure purchase decisions better understand what to look for from storage vendors that reduce power consumption and the total cost of ownership of enterprise storage. It is not a matter of if but a matter of when power efficiency considerations are going to start having a major impact on storage infrastructure purchases for enterprises that maintain a large amount of on-premises IT equipment. To be able to recognize and choose the most power-efficient storage systems requires an understanding of considerations beyond what many enterprise storage vendors like IBM are willing to openly discuss today.

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