

ESG SHOWCASE

Dell PowerScale Makes Its Performance Story Even Stronger

Date: October 2022 **Author:** Scott Sinclair, Senior Analyst

ABSTRACT: In the data-driven economy, file data can no longer simply be restricted to the spinning media of traditional IT. Unlocking the business potential in file data requires the performance of all-flash hardware. Given the massive capacities common to modern file environments, NVMe-powered storage becomes essential. With its newest OS and hardware, Dell PowerScale is delivering this transformational file storage infrastructure. As a result, organizations can power new workloads, save money, and strengthen their competitive positions.

Overview

The world is changing quickly, and companies need to be ready for what's next. Innovating with data has been playing an increasing role in determining business competitiveness. Digital information is a corporate asset as valuable as any other—in some cases, more so. Using digital information creatively, effectively, and efficiently can help a company differentiate itself in the marketplace. That achievement, in turn, reduces the company's risk profile and creates new revenue opportunities.

Simply put, in a modern data-driven economy, business success—or conversely, business struggles—can equate directly to how well a company maximizes the value of its data, especially its file data. This may be easier said than done because, according to ESG research, 46% of those surveyed believe their IT environment is more complex today than two years ago.¹

This complexity is fueled in part by the emergence of modern workloads (e.g., artificial intelligence, machine learning, deep learning, internet of things, and data lakes) that drastically alter the demands on the file storage environment, making it necessary for file storage infrastructures to evolve. According to ESG research, organizations are deriving more value, and incremental revenue, from their data as analytics workloads have become more pervasive on-premises and in the cloud.² Businesses today need low-latency, high-performance access to the vast bulk of their corporate file data, not just subsets of it. File storage no longer has to be just big; it must be big, fast, flexible, and simple to manage – while always keeping the data secure.

File storage no longer has to be just big; it must be *big, fast, and simple to manage.*

The requirements associated with massive capacity scaling plus high performance can strain IT budgets. It is a situation [Dell PowerScale](#) is now addressing with its family of powerful all-NVMe platforms and a new release of the PowerScale OneFS operating system.

¹ Source: ESG Complete Survey Results, [2022 Technology Spending Intentions Survey](#), November 2021.

² Source: ESG Research Report, [Data Storage Trends in an Increasingly Hybrid Cloud World](#), March 2020.


These solutions come with blazing performance and massive capacity to deliver high-performance access to files while minimizing the storage footprint and costs traditionally associated with an ultra-fast, high-capacity, business-critical file environment.

If You Want to Maximize the Value of Your Digital Assets, You Need High-performing File Storage

According to ESG research, file data now appears to be the leading driver of enterprise data growth. IT managers were asked which workloads they believed would be responsible for their organizations’ on-premises storage and public cloud storage growth for the following 24 months. Three of the top four most common responses were internet of things, big data/data lake repository for unstructured data, and AI/ML workloads.³

Figure 1. Top Four Workloads Responsible for Data Growth On-premises and in the Cloud

Which of the following applications/workloads do you believe will be responsible for your organization’s on-premises storage and public cloud infrastructure spending growth over the next 24 months? (Percent of respondents)



Source: Enterprise Strategy Group

Connections exist among this file growth, increases in file-related business value, and the success or failure of company-wide digital transformation programs. Consider how important business intelligence/analytics has become to the success of modern businesses. Analytics provides actionable insights to internal analysts and decision makers, leading to innovation with data in near-real time—a great example of how one workload will spearhead the need for IT to deploy high-performance file storage.

Other workloads requiring extreme performance and scalability, such as those associated with artificial intelligence, 8K media and entertainment, or financial analytics, have just as much file-related business value.

Modern Workload Demands Have Transformed File Storage Requirements

No company wants to be left behind, and no company wants to be limited by a traditional file storage system unable to keep pace with the demands of a digital business. But traditional file storage hardware and operating systems focus mostly

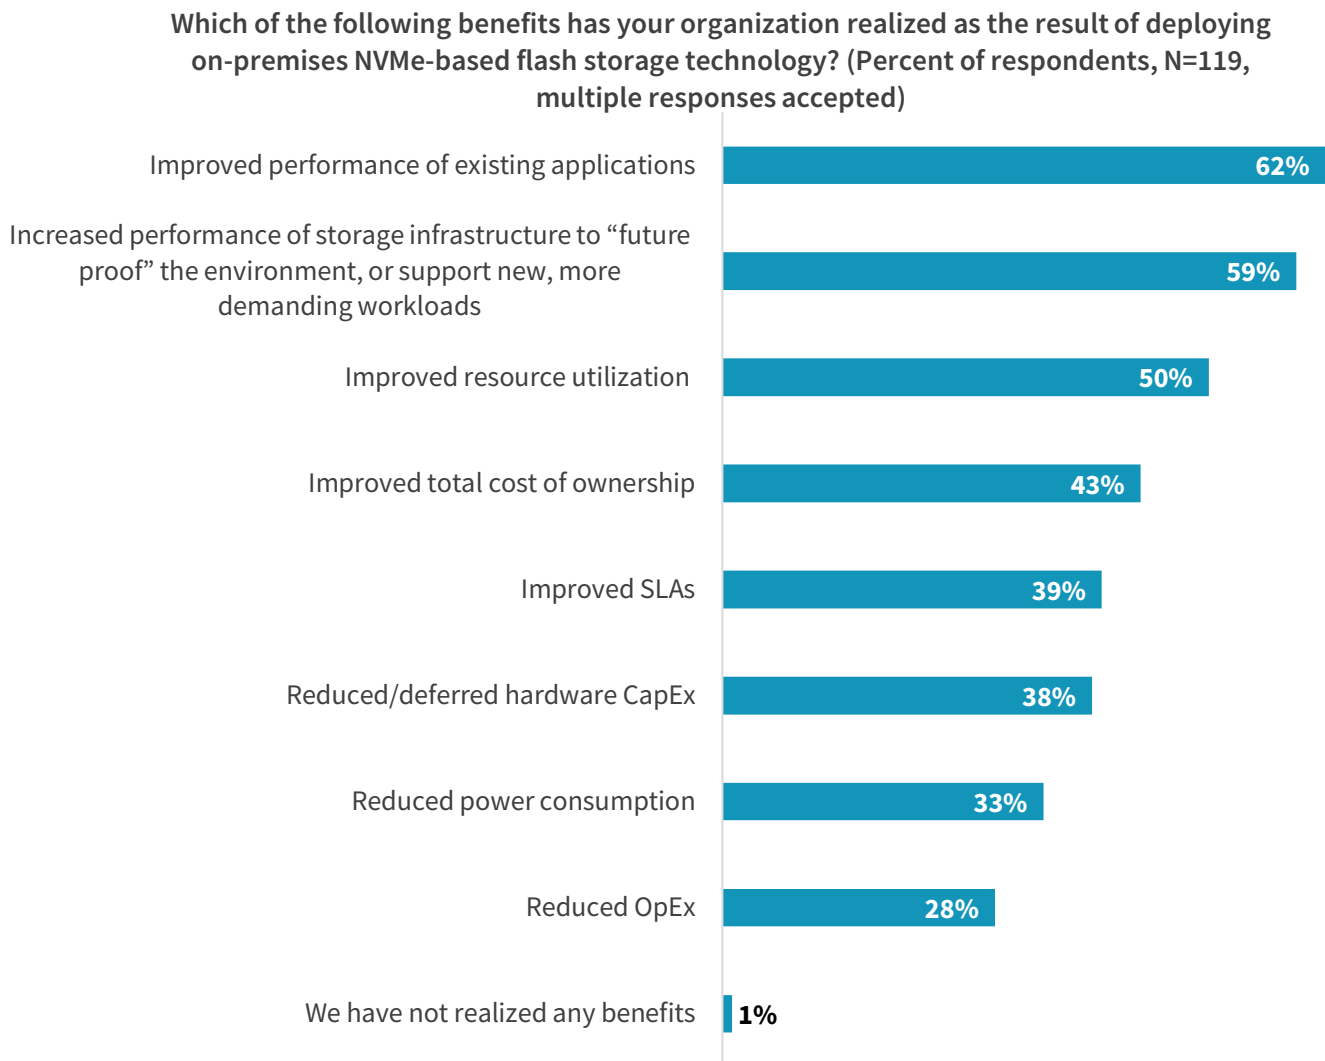
³ Ibid.

on delivering scale. The general assumption has been that only a small percentage of file data is going to be “hot” or active, with the bulk being cold and inactive. Expectations for traditional file storage are tied to the assumption that file accesses would be rare, and even when they did occur, they would predominantly be read-only operations. Often, performance is an afterthought.

Analytics has really upended that scenario. Insight is now perishable, so to speak. Business intelligence has a short shelf life and must be harnessed in real or near-real time to gain benefit. Therefore, ESG believes *file storage that supports analytics needs not just flash, but NVMe-powered flash.*

Flash storage was initially invented mainly to support OLTP environments that require speed, but today it is also transforming file environments. NVMe-based flash storage provides numerous benefits, including the ones shown in Figure 2.⁴

Figure 2. Benefits of NVMe-based Flash Storage Usage



Source: Enterprise Strategy Group

⁴ Source: ESG Survey Results, [2021 Data Infrastructure Trends](#), September 2021.

In addition to the realized benefits of improved application performance for existing application (62%) and increased performance to “future proof” the environment (59%), the efficiencies gained by moving from disk-based storage to flash have helped reduce costs as well. For example, nearly half of respondents identified improved TCO (43%). In addition, over a third identified reduced/deferred hardware capital expenditures (38%), and close to a third identified reduced operational expenses (28%).⁵ These benefits can be further improved if the NVMe-based flash storage infrastructure leverages inline data reduction to minimize the capacity required. Businesses need to innovate with data, and NVMe accelerates those efforts.

Transform File Storage with the PowerScale all-NVMe Family

A solution coming to the rescue is the Dell PowerScale all-NVMe scale-out NAS storage array with inline compression and deduplication technology, which is powered by the latest version of the OneFS operating system.

An Impressive Performance Story

The PowerScale OneFS operating system has been significantly enhanced to take advantage of the massive scalability and performance of PowerEdge-based platforms. For example, the 2U PowerScale F900 multi-CPU platform supports 24 drives and up to 100 GbE connectivity. The F900 cluster can start with a minimum of 3 nodes and can incrementally scale up to 252 nodes. The slightly smaller F600 has less capacity but delivers similar performance in a smaller 1U form factor. The F600 includes a choice of CPU and memory options, which enables it to deliver near F900 throughput to meet the needs of demanding workloads.

Both the F600 and F900 nodes can contain NVMe drives for accelerated performance. In 2022, Dell Technologies released 30 TB QLC-based drives that are available for the F600 and F900 nodes. The addition of QLC scaled the cluster capacity up by 2x without any loss in performance, according to information received from Dell. Increasing the cluster capacity up to 186 PB raw can significantly reduce the storage footprint and also the power and cooling requirements in most configurations. Environmental factors like sustainability are a key business driver for many enterprises who are looking at “greener data centers” and the new PowerScale nodes with QLC drives can be an integral part of the sustainability story.

PowerScale also released a new P100 performance accelerator that can benefit certain workloads that don’t need capacity but can benefit with higher throughput and lower latency. Adding P100 nodes to the cluster provides a lower cost way to scale the performance needs of demanding applications.

The PowerScale nodes and new OS seamlessly exist in the same cluster with previous generations of Isilon storage nodes, thus future proofing storage investments. The All-NVMe nodes can also seamlessly be added to clusters that include the newer PowerScale Hybrid (H700 and H7000) and Archive (A300 and A3000) nodes. The latest announced Performance accelerator (P100) delivers more performance without any storage to address certain key workloads.

Powered by the latest version of OneFS, the PowerScale scale-out NAS array is likely going to be close to ideal for supporting unstructured data workloads that demand extreme performance and efficiency. Notably, the operating system requires no configuration, is default-on, and provides continuous compression/decompression. Features like NFS over RDMA and NVIDIA GPUDirect support broaden the ability of the nodes to support performance-demanding workloads as evidenced by this report. A technical white paper from Dell, [Dell EMC PowerScale and NVIDIA GPUDirect](#), about PowerScale accelerated solutions with NFS over RDMA, shows linear scalability in the face of ever-changing analytics workloads.

A recent ESG brief, [The Data Center Lives On, but Its Future Will Be Very Different](#), highlights the need for scalable, modernized infrastructure going forward, with 63% of IT professionals anticipating that their organizations will be

⁵ Ibid.

operating six or more data centers within five years.⁶ Given these parameters, optimizing power, cooling, and floor space are important. With the massive growth in unstructured data that more than doubles every few years, there is an increasing need for organizations to reduce their data center power consumption. A QLC-based PowerScale offering can deliver increased capacity and performance while still being a cost-effective solution in demanding workload environments.

With the option of a cloud-like, as-a-service operating model for file services provided via APEX Data Storage Services for File data, Dell offers a simple option to reduce the upfront cost of storage infrastructure acquisition while offloading the operations to Dell. As a result, organizations achieve the improved simplicity and agility of cloud-like operations, maximizing their investments in data center infrastructure while maintaining the control associated with data center operations.

Speeds and feeds are one thing, but ultimately, it's the flexibility and enterprise grade features of OneFS-like data protection, data management, security, and performance management that make it simple to manage at any scale. The vast ecosystem of partners and system integrators enables customers to run a large number of workloads so they can innovate with their data. The key for PowerScale is to provide flexibility and choice to address the workloads of today and the unpredictable needs of tomorrow.

Use Case Example:

Algorithmic quant trading firms are in the midst of transitioning to support growing quant teams on increasingly larger data sets as they move their modeling from intra-day to multi-day trading. This, coupled with the exponential growth in daily transactions, means that quants can no longer store the active trade data sets in memory, and today's alternative solutions struggle to deliver the required performance at scale.

PowerScale nodes deliver the high performance and extreme concurrency at scale from the edge to the entire data set. This shortens model development time with faster analysis on larger, multi-day trading data and extreme performance for smaller real-time data. Combined with its enterprise features, this enables a simple, efficient solution, which accelerates cycles of learning by bridging historical and real-time databases, while conforming to regulatory standards and guaranteeing enterprise data protection. PowerScale also offers a smaller footprint and lower cost of entry while delivering the powerful performance characteristics required.

The Bigger Truth

By offering the PowerScale solutions with massive performance and capacity characteristics for demanding workloads, Dell is actively and directly addressing the challenges of modern file storage demands. With the PowerScale family, Dell is focusing not only on reducing the hardware costs resulting from rapid file data growth, but also on maximizing the low-latency performance potential. The ecosystem of partners supported by the OneFS operating systems provides peace of mind while maximizing the business benefits to users.

⁶ Source: ESG Brief, [The Data Center Lives On, but Its Future Will Be Very Different](#), September 2022.



All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.



www.esg-global.com



contact@esg-global.com



508.482.0188