White Paper

Meeting the New Unstructured Storage Requirements for Digitally Transforming Enterprises

Sponsored by: Dell Technologies and Intel

Eric Burgener
June 2021

IDC OPINION

Digital transformation (DX) is underway among 90% of enterprises globally. As information technology (IT) organizations evolve to support more data-centric business models, infrastructure modernization will continue throughout enterprises. Most of these projects are driven by the need to support new workload requirements imposed by numerous next-generation applications (NGAs) being deployed as part of DX. The majority of data being collected to drive better business insights in the new digital era is unstructured. In addition, the focus on making data-driven decisions is imposing new performance, availability, cloud integration, flexibility, and ease-of-use requirements on the scale-out architectures that are being used to capture, store, protect, and analyze this unstructured data.

To implement effective unstructured data management strategies, CIOs must evaluate the applicability of many new technologies. Those technologies include artificial intelligence (AI), accelerated compute, NVMe and related solid state storage technologies, and unified namespaces that span both file- and object-based storage on premises and in the public cloud. They also include the infusion of enterprise-class data management features into massively scalable distributed architectures. In an attempt to simplify data pipelines for data-driven NGAs, CIOs are looking to reduce the number of independent storage silos by consolidating workloads onto fewer high-density platforms. This drives new requirements for supporting multiple access methods, storage tiering, flexible quality-of-service (QoS) controls, and automation that functions across bare metal, virtualized, and containerized environments.

After reviewing recent IDC primary research exploring enterprise customers' challenges and requirements for unstructured data management, this white paper will cover a brief overview of Dell EMC's unstructured data storage portfolio. The vendor offers a comprehensive ecosystem for managing file, object, and streaming data storage requirements, complete with AI-driven infrastructure monitoring, predictive analytics, and data life-cycle management tools. Dell EMC offers its customers the blueprint and solutions required to deploy an agile unstructured data management strategy that will help them achieve their DX objectives across edge, core, and cloud-based infrastructure.

IN THIS WHITE PAPER

In 2021, IDC conducted primary research specifically focused on enterprise unstructured storage requirements (IDC's *Unstructured Data Storage [UDS] Survey*). The survey targeted storage purchase decision makers and influencers across industries in North America (the United States and Canada)

with a sample size of 406. Survey respondents identified features and capabilities needed for their unstructured storage platforms going forward as well as workload affinities and preferences for different deployment models (on-premises and public cloud–based file storage, on-premises and public cloud–based object storage). This white paper summarizes those survey findings, suggesting specific IT management considerations for accommodating NGA deployment as a result of DX.

SITUATION OVERVIEW

Over 90% of enterprises are undergoing DX and evolving their data-centric decision making and business models. More data than ever before is being captured, stored, protected, and analyzed. Over the next five years, 80% of that data will be unstructured (i.e., file- and/or object-based) data. Primary research conducted by IDC in late 2020 indicates that almost 70% of enterprises will modernize their server, storage, and/or data protection infrastructure within the next two years. These technology refresh projects are primarily driven by the requirements of new workloads that enterprises are deploying as a result of DX — many of which leverage AI-driven big data analytics workloads using unstructured data.

These next-generation workloads demand a set of capabilities that legacy storage platforms are not well equipped to address. The combination of massive volumes of unstructured data, the increasing importance of this data in day-to-day decision making, and the need for faster analysis to drive real-time business decisions is resulting in a renaissance among unstructured storage platforms.

Overall, enterprises have an average of 71PB of file storage in on-premises locations and over 91PB of it in public cloud locations. On-premises file storage was expected to grow at 46% per year over the next five years, while public cloud file storage was expected to grow even faster (at 53% per year). With respect to object storage, enterprises had roughly the same amount (71.7PB) on premises but about 18% less in the public cloud (77.1PB). Growth rates for data in the public cloud, fueled by both file and object storage growth, were higher than growth rates for on-premises storage.

As shown in Figure 1, the top drivers of this data growth on premises are DX-related factors such as new AI-driven projects, a general increase in digital data, and increased amounts of data (either due to higher resolution or more frequent collection) generated by existing sources. Edge computing deployments, where additional data can be collected to increase overall visibility into the business, also contribute, driving an increased data protection load as well as the need for more efficient, higher-speed networking technologies. Edge locations are on the rise, with only 35% of enterprises having fewer than 10 of them. 34% of enterprises have between 25 and 50 edge locations, and 3% of enterprises have more than 50. This underscores the importance of having cost-effective, compact, and efficient edge unstructured storage solutions to help collect, store, protect, and analyze this data.

FIGURE 1

Top Drivers of On-Premises Data Growth

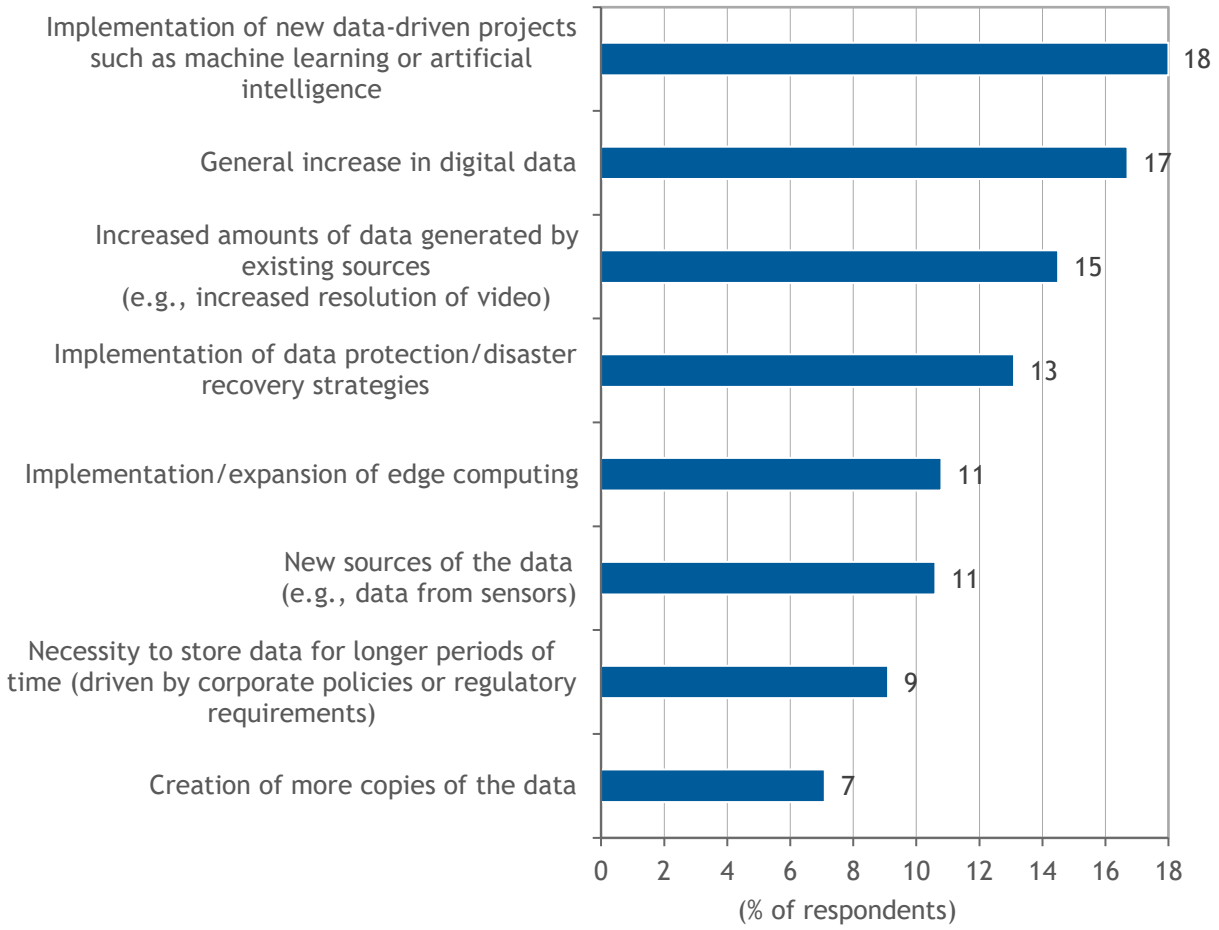

Source: IDC, 2021

Many organizations have deployed tiered storage solutions, allowing data to be kept in the appropriate tier based on performance, availability, capacity, and cost requirements. Those solutions may support multiple media types within a single system, using intelligent data placement algorithms to determine whether data is stored on all-flash, hybrid, or hard disk drive (HDD)-only storage tiers. With the increasing size of data sets and the need to retain data for longer periods to meet compliance and other requirements, the ability to tier data to S3 targets (whether those are on-premises or public cloud-based object storage) has become fundamental for unstructured storage solutions.

With roughly 35% of all corporate data not used within the past six months, archiving is an important strategy for increasing the efficiency of storage infrastructure associated with more active data and for lowering overall costs. Enterprises leverage both on- and off-premises options for this task, with 52% of them tiering older data to on-premises archives and 47% of them tiering to a public cloud archive (keeping more active data on premises). The broad adoption of archiving strategies underlines the need for data classification tools that drive intelligent data placement based on IT-defined policies.

Although traditionally big data analytics environments have used a storage infrastructure with a distributed, high-performance file system on the front end and an S3 target on the back end, intelligent intrasystem tiering capabilities, increased system scalability, and a range of new storage media types have made it increasingly possible to cost effectively create a single system that can store both active and archive data. As an example, cost-effective and highly scalable object storage that includes a high-performance persistent storage tier using all-flash media is being deployed by 12% of enterprises today, and an additional 8% plan to deploy it within the next 12 months.

More enterprises are looking to reduce the number of unstructured storage silos they must maintain, primarily to simplify administration, enable better data sharing across workloads using different access methods, and lower costs. Denser workload consolidation begs the performance question, but newer technologies like NVMe, storage-class memory (i.e., Intel Optane–based storage devices), and AI-assisted performance management go a long way toward addressing these concerns. Support for different access methods is critical. For analytics data sets that have historically been more batch oriented, NFS, SMB, S3, and HDFS are important. For those enterprises moving in the direction of containerized, cloud-native workloads, two interfaces will be important: the Container Storage Interface (CSI) for block- and file-based storage and the Container Object Storage Interface (COSI) for object-based storage (a standard that is still under development). Simple, nondisruptive scalability into the tens of petabytes range is a minimum requirement for customers looking to consolidate workloads — each of which may be experiencing high data growth — onto fewer platforms.

Many of the NGAs being deployed as part of DX are more critical to day-to-day business operations than the workloads of scale-out unstructured storage platforms in the past. This has significantly increased both the availability and the recovery requirements for these platforms. For customers looking to refresh on-premises file storage, availability concerns dominate, outstripping even performance, and providing the high availability needed for mission-critical workloads can often present challenges to legacy storage infrastructure. Figure 2 shows that the top concerns — nondisruptive performance and capacity scaling, high availability, and business continuity — are all related to minimizing the opportunities for downtime on these systems. Hardware features like redundant, hot-pluggable, field-replaceable components and software features like multiple data protection options (RAID or erasure coding), snapshots, replication, and integration with enterprise backup solutions are becoming increasingly important for unstructured storage platforms in digitally transforming enterprises.

FIGURE 2

On-Premises File-Based Storage Concerns

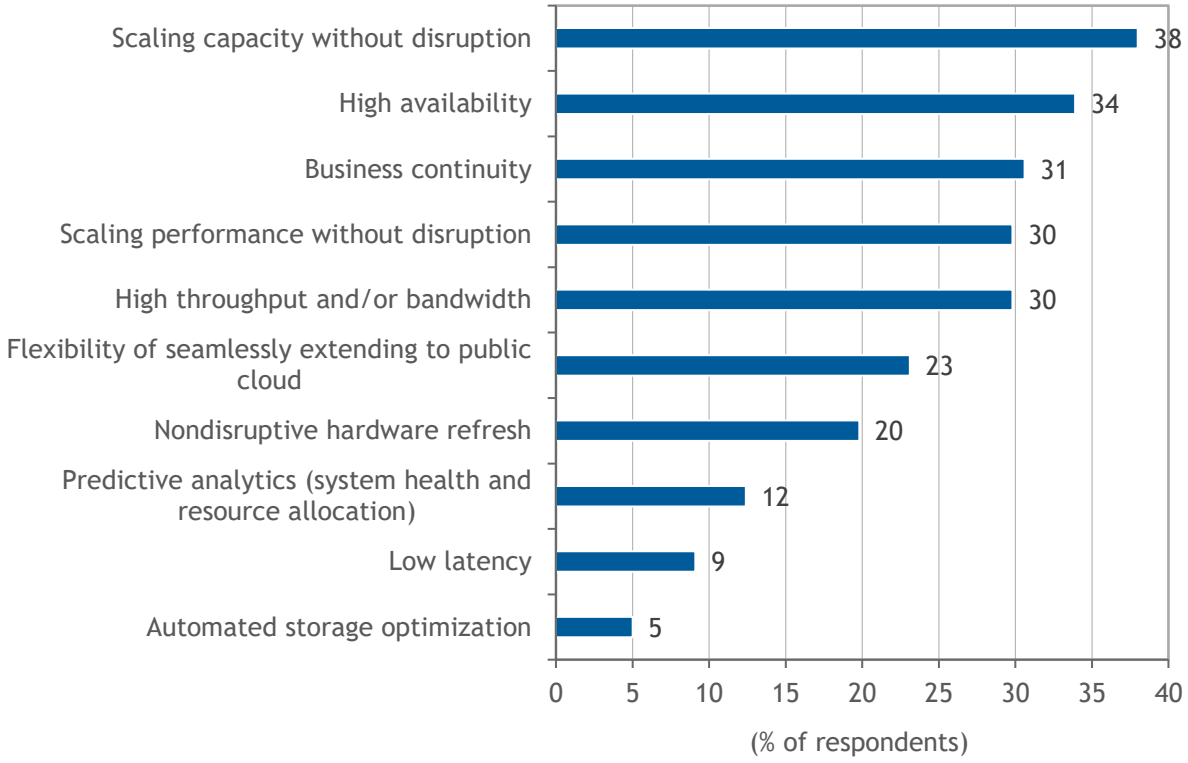

Source: IDC, 2021

With the increasing pace of business, IT agility is a critical enabler of financial success. Along with lower cost to support certain types of workloads, agility was one of the initial advantages that public cloud-based services offered to enable quick response to new and/or changing business conditions. The ability to outsource IT infrastructure management, easily leverage newer technologies that could not yet be cost effectively deployed in-house, and shift IT spend toward operational (rather than capital) expenditures was attractive. Most vendors today offer the option to buy in-house storage infrastructure using opex models, removing that as a feature unique to cloud-based offerings, and for most of the larger storage vendors, the subscription-based services business is growing rapidly. Most customers appreciate the option to tie costs more closely with actual usage.

Key Features of Unstructured Storage Solutions

With the world of unstructured storage management changing rapidly due to DX, we asked our survey respondents to identify the most important attributes of file- and object-based storage. For file-based storage, the top workloads on premises (content and security applications and unstructured big data analytics) were somewhat different from those in the public cloud (archiving, unstructured big data analytics, web serving, and backup/disaster recovery [DR]). Many different file system features were considered important, with the top 9 all being identified as critical for between 24% and 21% of respondents (see Figure 3).

FIGURE 3

Important Features of File-Based Storage

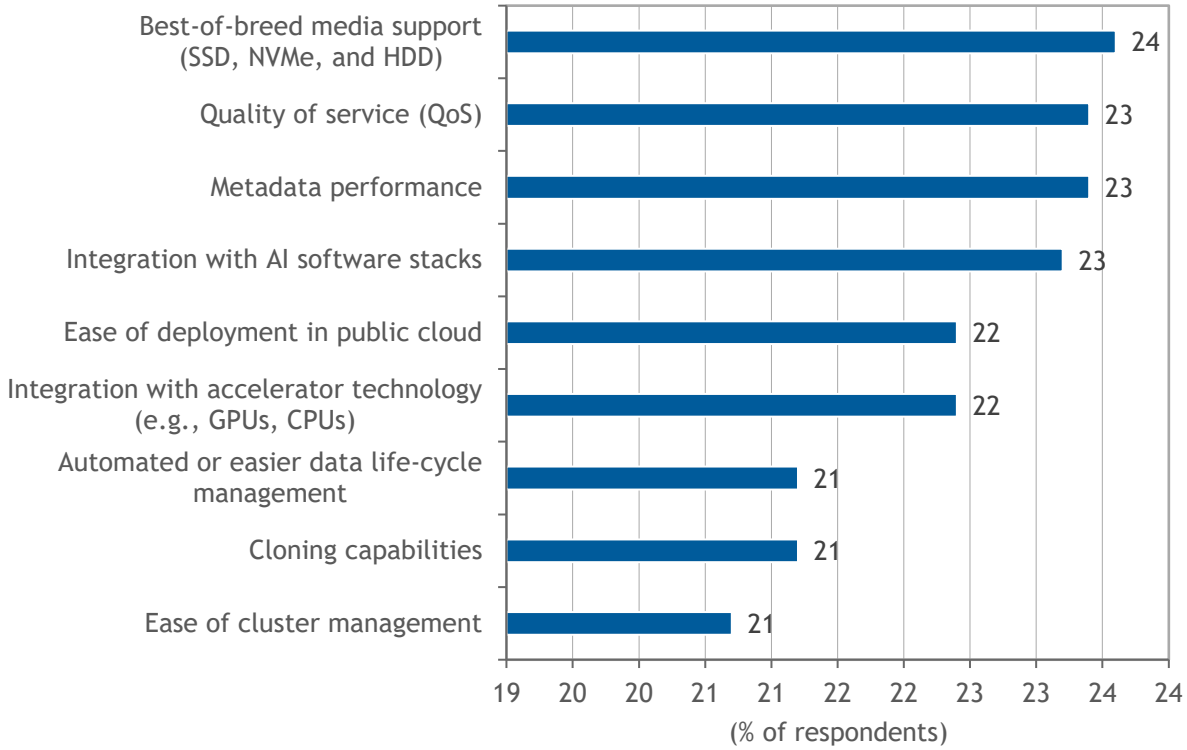

Source: IDC, 2021

The flexibility of best-of-breed media support won out as the top requirement by a small margin. The ability to support a variety of different media types, and in particular new media types as they become available, enables administrators to create tiers of storage within a single system that achieve the optimum trade-offs between performance, capacity, and cost based on their workload mix. Coming in at number 2, quality of service is important for enterprises that must meet specific performance objectives even as they add more workloads to a single platform. Rounding out the top 4 were metadata performance and integration with AI software stacks, features that are critical to performance in many AI-driven workloads. Note that the number 4, integration with AI software stacks, was supported by 22% of respondents identifying integration with accelerated compute technology as important as well (coming in at number 6).

Ease-of-use issues appeared three times in the top 9 features identified. Most enterprises are already operating in hybrid cloud environments, and the ability to deploy the same file-based offerings either on in-house or public cloud infrastructure can help ease management issues and better enable optimal workload placement as applications evolve over time. Easier cluster management in general is important, particularly as the responsibilities for a single storage administrator increase over time. Enterprises are clearly investing in more automation and orchestration tools like Ansible and Kubernetes for easier data life-cycle management.

For object storage, the top workloads it was used for, both on premises and in the public cloud, were archiving and content applications. Usage, however, is changing with the advent of NGAs. Given the value of very large data sets to many AI-driven workloads, enterprises are looking for features that significantly increase the performance of object stores so they can use them for more performance-sensitive workloads. Top object storage features our survey respondents identified going forward included integration with accelerated compute technologies, QoS, integration with AI software stacks, and best-of-breed media support (in particular supporting the definition of higher-performance storage tiers leveraging flash media) (see Figure 4).

FIGURE 4

Important Features of Object-Based Storage

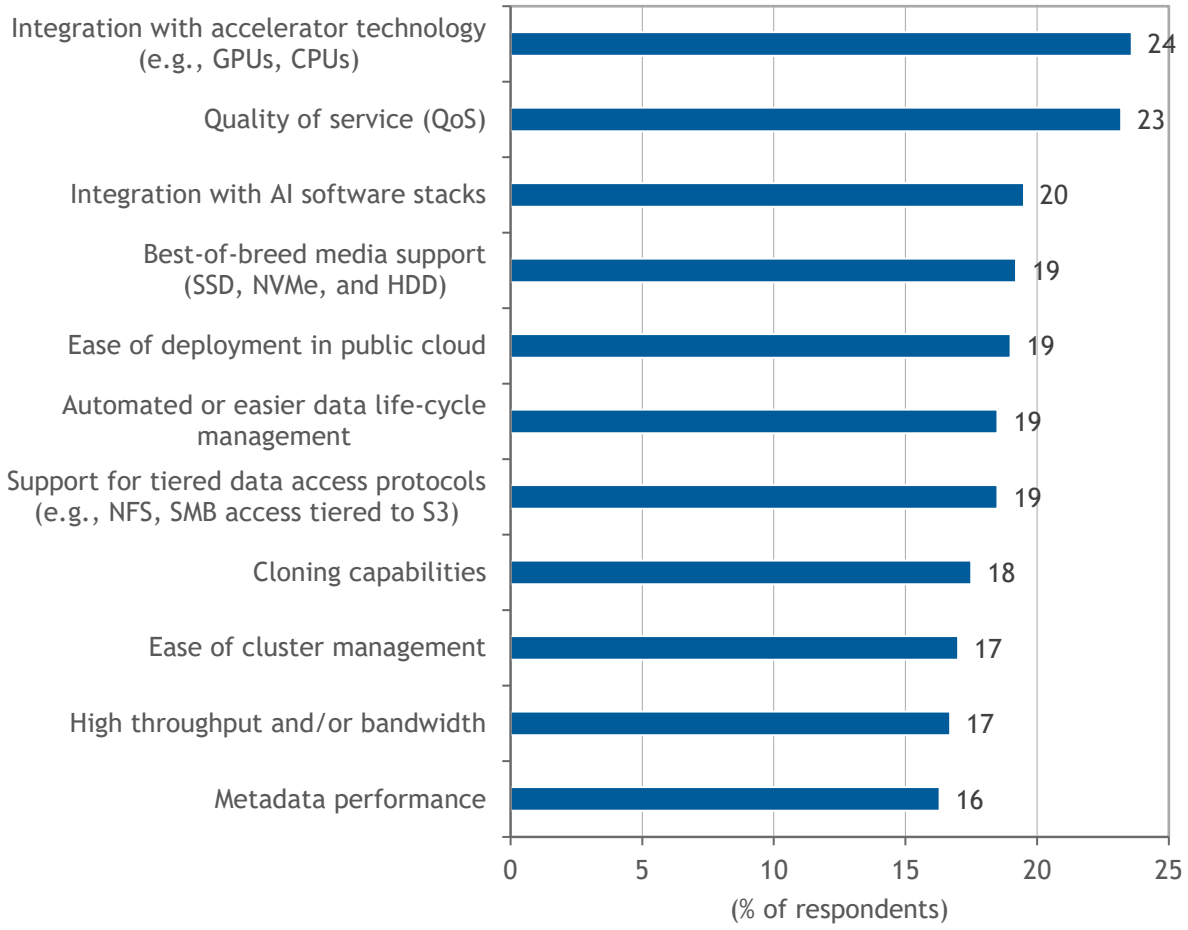

Source: IDC, 2021

Ease-of-use features – easy public cloud deployment and automated data life-cycle management – came in at number 5 and number 6, respectively, with 19% of respondents identifying them as important. The need for multiple access methods was mentioned previously in the paper, but here it shows up as the number 7 important feature with almost 19% of respondents also calling it out.

The Increasing Importance of Artificial Intelligence-Driven Workloads

The advent of accelerated compute, higher CPU core counts, and higher-performance storage technologies has accelerated the use of AI in business. From big data analytics to machine learning (ML) and deep learning workloads, enterprises are driving differentiating business value using AI technologies. For many of them, having more data available for analyses drives better outcomes. While AI workloads may need high ingest (write) throughput like some more traditional performance-intensive computing (i.e., high-performance computing [HPC]) workloads, they also often need to perform many small, random reads with low latency (metadata-intensive operations for which traditional scale-out file systems have not been optimized). This is driving demand for a new type of scale-out storage platform that can cost effectively deliver across a variety of different I/O profiles, not solely the "high sequential throughput against large files" that had been so typical of performance-intensive computing.

We asked survey respondents about their AI infrastructure. 29% of respondents were already running real-time, mission-critical AI workloads, with an additional 12% planning to add at least one of these NGAs within the next 24 months. Security and performance topped the wish lists for AI infrastructure with 43% and 42% (respectively) of respondents identifying them as the leading requirements. For storage, desired security features included flexible encryption that could be configured for data either in flight and/or at rest, role-based access control for data access, secure snapshots with retention schedules, and write once, read many (WORM) capabilities for archives. Storage performance technologies of interest included NVMe, RDMA, flash, and support for accelerated compute (e.g., GPU-based systems).

Respondents identified several barriers to successful AI adoption as well. The technology's relative immaturity was a concern, with 34% of enterprises calling this out as the number 1 issue. It is interesting to note that "proven technology" was identified by 33% of survey respondents as a requirement for AI infrastructure. Other barriers included a lack of executive buy-in (with 32% identifying that as a key issue), a lack of skilled talent (32% as well), and a lack of data (31%). From the interest in the use of AI in business, it is clear that it will eventually become a mainstream IT infrastructure technology, but it is also true that it is still relatively immature.

Although our survey did not specifically ask about the availability of converged infrastructure and reference architectures, it is clear from other AI research IDC has performed over the past 12 months that commercial customers like these options very much. Reference architectures address the AI infrastructure definition and cross-product validation issues, but converged infrastructure offerings go farther in offering a single SKU for easy purchasing and a single point of support contact as well.

The Dell Technologies Unstructured Storage Portfolio

Dell Technologies is a $92 billion multinational technology company whose portfolio includes a wide range of both personal and enterprise IT infrastructure. With its unstructured data storage offerings, Dell EMC (the vendor's enterprise storage business unit) has held the number 1 market share by revenue position for a number of years, including most recently in 2020. Throughout this period of hegemony, the vendor has driven innovation in its file- and object-based storage platforms for both on-premises and cloud deployments that keeps pace with evolving unstructured storage requirements in the era of DX.

In looking across the unstructured storage market, Dell EMC sees three types of unstructured data. File-based storage is based on a folder hierarchy and generally used to store audio, video, and photos as well as genomics and analytics data. Object-based storage utilizes a single address space, is optimized for massive capacity, and is used to store active and cold archives, data used by modern cloud applications, and increasingly backup and disaster recovery data. Streaming data is data that is constantly in motion, used with internet of things (IoT) and data-in-motion analytics. To address the requirements of these three types of unstructured data, the vendor's broad portfolio includes file-based storage, object-based storage, streaming data platforms, and unstructured data set management tools.

File-Based Storage (On Premises and in the Cloud)

The Dell EMC PowerScale/Isilon platform includes a family of systems that run OneFS, the vendor's scale-out NAS software platform. Available in NVMe-based all-flash and SAS-based hybrid and archive appliances, PowerScale can start with a cost-effective three-node cluster for edge and distributed locations and scale all the way up to a 252-node cluster that provides up to 15.75 million IOPS, 945GB/sec of aggregate throughput, and 30.96PB of raw storage per cluster (using NVMe-based nodes). Hybrid and archive nodes can be mixed and matched in a single cluster to provide the flexibility to efficiently hit capacity, cost, and performance goals, and each node type supports multiple drive options for additional configuration flexibility.

The OneFS operating environment provides many enterprise-class storage management capabilities along with its high-end scalability. With inline compression and deduplication, 256-bit encryption, erasure coding, snapshots and replication, storage tiering (SmartPools and CloudPools), quotas, retention protection, and file immutability (write once, read many), it provides a highly available and resilient parallel file system platform. The OneFS AutoBalance function distributes data across nodes for high performance and availability and supports non-disruptive expansion, data migration, or application logic reconfiguration. PowerScale's broad multiprotocol support, which includes NFS, SMB, S3, HTTP, FTP, HDFS, and IPv4 and v6, enables all data to be simultaneously read and written through any protocol, providing highly efficient data sharing without requiring data migration. PowerScale systems are available as appliances, built on top of Dell's PowerEdge server technology, or in the future as software only (which enables deployment in public cloud-based environments), giving customers the flexibility to support any unstructured data workload for any user, anywhere.

Organizations that enjoy PowerScale on premises have the freedom to choose the most suitable cloud or multiple clouds to drive the best business outcome. They can combine the massive scalability, performance, and data services that they know in their datacenters with the economics and simplicity of public cloud. Dell Technologies offers cloud solutions that are directly connected to multiple clouds, allowing customers to switch between them without moving or copying the data in order to benefit from the best compute and other native cloud services and natively integrated solutions that can be acquired and managed on the cloud provider's marketplace. In any case, the applications don't have to be refactored, allowing organizations to accelerate their migration to the cloud at reduced costs and lower risk.

PowerScale/Isilon systems are broadly used across many vertical markets to support a variety of workloads including media and entertainment, healthcare (e.g., DNA sequencing), research, technology, telecommunications, autonomous driving, agriculture, public safety (e.g., Smart Cities), general-purpose file sharing, and a variety of AI-based, performance-intensive (i.e., high-performance computing) and other technical computing workloads. Many of these workloads are broadly served in the cloud with the help of innovative native services offered by various cloud providers that can be bundled with the PowerScale cloud solutions.

Object-Based Storage

The Dell EMC ECS platform includes software and a family of systems that provide enterprise object storage with scalability into the multi-exabyte range. Available in NVMe-based all-flash as well as HDD-based appliances (which support optional SSDs for metadata caching), ECS can start with a five-node cluster and scale indefinitely (there is no upper limit on the number of nodes per cluster). ECS provides a single global namespace across geographically distributed environments and delivers high availability and resiliency features to deliver a true enterprise-class solution. Deployable as a turnkey appliance or in a software-defined model, ECS delivers rich S3 compatibility on a globally distributed architecture and is purpose built to store unstructured data at public cloud scale with the reliability and control of a private cloud. ECS also supports multiple protocols — NFS, CAS, SMB, CIFS, and HDFS — and allows for interoperability between certain protocols.

Dell EMC ECS appliances are built on top of Dell PowerEdge server hardware. Different appliance models are targeted at different use cases, including the low-end EX300 HDD-based node starting at 60TB for smaller environments to the extremely dense EX3000 (which can accommodate up to 11.5PB per 42U rack) and the high-performance NVMe-based EXF900 (which can scale up to 2.9PB per rack). The EX300 features multiple 10GbE network connections, while the larger models support 25GbE connections.

Traditionally, Dell EMC ECS has been used as a massively scalable data lake behind various parallel file systems as a cost-effective S3 tier for colder data, a content repository, and a long-term data retention platform (that meets compliance and regulatory requirements), as well as for archive, backup, and disaster recovery data. With the availability of an all NVMe-based version of ECS (as well as the ability to insert SSDs in any HDD-based node for metadata caching), the platform can be used for the emerging set of performance-oriented big data analytics workloads as well as AI, machine learning, IoT, and cloud-native applications.

Streaming Data Platform

The Dell EMC Streaming Data Platform (SDP) is an enterprise-ready software platform that supports a hybrid multicloud ecosystem for unstructured data. SDP defines a reference architecture, using Dell EMC storage products as well as third-party products, that covers the entire data pipeline from ingest through to storage, analysis, and archiving. Data collection devices such as sensors, cameras, medical devices, and autonomous vehicles (among many others) capture and relay data through various wireless network gateways to edge storage, where it can be processed and/or forwarded onto core Dell storage platforms (including various Dell EMC primary storage products as well as Dell multicloud offerings) for analysis, integration, and alerting. From there, data can be tiered to PowerScale, ECS, and/or private or public cloud (AWS, Azure, and Google Cloud) platforms for use in AI/ML model building, analytics, and management dashboards.

In general, the release of reference architectures has been extremely helpful to enterprises that are deploying new unstructured storage platforms and using them for more performance- and availability-sensitive workloads than scale-out architectures have been used for in the past. Dell EMC's SDP is built around Pravega, an open source platform, and provides a blueprint for customers looking to unify their data pipeline for next-generation AI and big data analytics workloads. Like most reference architectures, SDP turns creating the right storage platform for NGAs from a "science project" into a deterministic implementation that is easier to buy, is faster to install and configure, and deploys more reliably.

Infrastructure Visibility Under the New Data-Centric Business Model

AI is not just being used at the application level. Many IT infrastructure providers have implemented their own AI-driven SaaS applications that comprehensively monitor all aspects of storage and other infrastructure components, using ML to help pre-validate upgrades, identify anomalies, drive system optimization, leverage predictive analytics to de-risk system operation, and troubleshoot faster and more comprehensively (when that is required). Dell EMC bundles in CloudIQ, a cloud-based software platform that provides ML-driven proactive monitoring and predictive analytics to reduce risk, better inform performance and capacity planning, and improve administrative productivity for core, edge, and cloud deployments, at no extra charge with all Dell EMC storage purchases. CloudIQ supports a broad range of Dell Technologies products, including storage (ECS, PowerScale, PowerMax, PowerStore, PowerVault, Unity XT, XtremIO, and SC Series), data protection (PowerProtect DD and PowerProtect Data Manager), converged and hyperconverged infrastructure (VxBlock and VxRail), and SAN switching (Connectrix).

Successfully implementing the unstructured data pipeline to service the variety of NGAs that many enterprises are deploying is not just about deploying the right storage infrastructure but also about getting a handle on the data assets. DataIQ delivers a single-pane-of-glass view of all file and object data across PowerScale, ECS, third-party storage, and the cloud with data tagging and precision data mover capabilities, high-speed scan, search and indexing across heterogeneous systems, and reporting on data usage, access patterns, and performance bottlenecks. Because of the criticality of data set management in the new digitally transformed business climate, DataIQ is bundled in with all Dell EMC unstructured data storage platforms at no additional charge, giving users the visibility they need to rapidly discover, understand, and act on novel data insights that drive better business decisions.

Flexible Consumption Models

There are five enterprise storage consumption models — appliances, software only, converged infrastructure, hyperconverged infrastructure, and services (i.e., cloud) based. Each consumption model offers different benefits and is targeted at different customer requirements. Not all IT infrastructure vendors offer their products across all models, but with its unstructured storage portfolio, Dell Technologies covers four of them and in the future will add software-only options. In other research performed by IDC in 2020 (IDC's *Infrastructure Portfolio Survey*), Dell customers identified the maturity of infrastructure solutions and the availability of multiple consumption models from Dell as two of the top reasons to purchase solutions from the vendor. It was noted previously that both PowerScale and ECS are available through multiple consumption models, and the software-defined nature of both of those platforms lends itself very well to making the products available with that level of flexibility.

Having been exposed to the "cloud experience" through various public cloud services, many enterprises are looking to evolve the in-house IT infrastructure "customer experience" more in that direction. Vendors have responded, providing better technology refresh models and subscription-based pricing options for in-house infrastructure purchases, and Dell Technologies has been no exception. The vendor has long offered the ability to purchase Dell EMC storage infrastructure for in-house deployment on an operational expenditure (i.e., opex) basis through the Dell Technologies on Demand program. In May 2021, the vendor introduced Project APEX, a comprehensive strategy for delivering a radically simplified as-a-service and cloud experience, based on Dell EMC's "Power" storage portfolio of products. The program provides fully managed infrastructure options, starting with storage services offerings that are available today, and unified management (through Dell Technologies Cloud Console) across private, public, and edge cloud operating models. Project APEX gives customers the option for "pay as you go," "pay as you use," or subscription-based pricing with elastic performance and capacity that can be scaled up or down.

CHALLENGES/OPPORTUNITIES

End users will be navigating many challenges as they refresh IT infrastructure to best accommodate the requirements of NGAs deployed as part of DX. AI-driven workloads will make up a major portion of those new applications, and across the AI data pipeline (ingest, label, train, validate, infer, and retain), there are many different storage requirements that in the past have required different storage silos (resulting in high complexity and cost). Consolidating the data pipelines in the new, more big data analytics–oriented business environments holds out the possibilities of better efficiencies and significant cost savings, but storage vendors need to understand the I/O requirements of each stage and explain clearly how their system architectures can meet those widely varying needs.

This challenge presents an opportunity for Dell EMC, however, with its unstructured data storage portfolio and related reference architectures. Enhancements over the past two years in leveraging newer NVMe-based technologies, supporting additional access methods, enabling better cloud integration, providing more consumption model options, moving to more software-defined architectures, and more fully instrumenting both systems and data for monitoring, management, and usage with AI-driven platforms meet these new requirements. All of this functionality provides good opportunities for competitive differentiation for Dell EMC in the unstructured data storage market.

CONCLUSION

DX drives enterprises toward a much more data-centric business model, and unstructured data will be the primary fuel for that engine. Given the size of the data sets, scale-out architectures provide the easy scalability needed in these environments, but NGA requirements demand many new capabilities that legacy scale-out architectures will struggle to deliver. This is driving a renaissance not only in unstructured storage platforms but also in how enterprises manage the unstructured data life cycle.

Availability, performance, cloud integration, flexibility, and ease-of-use features are the primary areas our survey respondents identified as key requirements as they purchase new or refresh older unstructured storage platforms as part of DX. AI-driven workloads will make up a good percentage of the NGAs deployed as part of this transition, and enterprises identified a particular set of concerns around AI adoption including technology maturity, lack of executive buy-in, an inability to find skilled talent, and a lack of data. An increasing number of AI workloads will be real time and mission critical, with 29% of survey respondents already running at least one workload of this type and an additional 12% planning to deploy one in the next 24 months. These types of workloads have even higher performance and availability requirements than more traditional batch-oriented analytics, and this is resulting in a rapid penetration of solid state technologies into file- and object-based storage platforms.

With its unstructured data storage portfolio — which includes Dell EMC PowerScale/Isilon (at the edge, in the core datacenter, and in the cloud), Dell EMC ECS, the Streaming Data Platform, and AI-driven tools for both infrastructure and data management (CloudIQ and DataIQ) — Dell provides the capabilities enterprises need to implement an effective and efficient unstructured data management strategy.

MESSAGE FROM THE SPONSOR

Intel Partnership

For more information on Dell Technologies unstructured data products of PowerScale, ECS and Streaming Data Platform, go here: www.delltechnologies.com/unstructureddata

Intel has partnered with Dell Technologies to optimize Intel® Xeon® Scalable CPUs, Intel® Optane™ SSDs, and Intel® 3D NAND SSDs for storage infrastructure and data services, delivering predictable performance and a seamless storage experience. Intel drives the industry forward in developing new technology standards and contributing to open standard projects.

Intel® 3D NAND SSDs work with Intel Xeon processors to meet today's increasingly demanding service levels and support broader cloud workloads, while reducing storage costs. Packed with a deep feature set, these cloud-20 inspired SSDs are optimized for cloud infrastructures, offering outstanding quality, reliability, advanced manageability and serviceability to minimize service disruptions.

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.