

# IDC MarketScape: Worldwide Distributed Scale-Out File System 2022 Vendor Assessment

Eric Burgener

## IDC MARKETSCAPE FIGURE

FIGURE 1

### IDC MarketScape Worldwide Distributed Scale-Out File System Vendor Assessment



Source: IDC, 2022

Please see the Appendix for detailed methodology, market definition, and scoring criteria.

## IDC OPINION

---

Over the next five years, scale-out file systems will be widely deployed by enterprises looking to consolidate file-based workloads, improve file-based infrastructure efficiencies, and handle many of the performance and scalability requirements of modernized applications that are very data intensive. All of the products evaluated here will be able to do that very well for most enterprises, although there are some differences in top-end performance and scalability and ease of use between offerings – that is why Figure 1 has many of the vendors clustered closely together. What the reader should note, however, is that there can be significant differences between vendors in their architectures, product strategies, areas of focus, and software-defined flexibility that should be evaluated as purchase decisions are made.

The "Advice for Technology Buyers" section is probably the most important section to read for those who will be involved in making a purchase decision. This section introduces a number of strategic questions enterprises should ask themselves when determining what is most important in selecting a scale-out file system offering. As an example, all evaluated products can support a 1PB file system, but what each system looks like, how easy it is to manage and upgrade, how much it costs and, in general, how it gets there can be very different. There is no "best" offering in this market, but there are certain products that are better suited for certain workloads and will cater better to certain objectives like top-end performance and scalability, ease of use and management, lower energy and floorspace consumption, hybrid cloud capabilities, and how different access methods are supported.

Enterprises can expect a lot more innovation to occur in the scale-out file market going forward, driven primarily by the fact that 80% of the data that will be created over the next five years will be file and/or object based. If enterprises just need to simplify basic file sharing (home directories, etc.), there are a lot of very viable options (some of which are mentioned in the "Vendors to Watch" section). Modernized applications, particularly those using artificial intelligence (AI) or those which are very data intensive, will have additional demands that may not be well met by the simpler products, and that's where enterprises will need to turn to true distributed scale-out file system platforms.

## IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

---

This IDC study assesses the capabilities and business strategies of popular suppliers in the distributed scale-out file-based storage market segment. For a complete definition of distributed scale-out file systems (and a discussion of the new file-based storage taxonomy that IDC introduced in July 2021), see *Reclassifying File Storage – A New Approach for the Future of Digital Infrastructure* (IDC #US48051221, July 2021). This evaluation is based on a comprehensive framework and a set of parameters that gauge the success of a supplier in delivering a scale-out file-based storage solution to the enterprise market.

To be evaluated in this study, a vendor needs to have a scale-out file-based storage platform:

- **That conforms to IDC's taxonomy.** According to *Reclassifying File Storage – A New Approach for the Future of Digital Infrastructure* (IDC #US48051221, July 2021), assessed products need to meet the definition of a distributed scale-out file system platform or a clustered scale-up file system that is sold primarily against distributed scale-out file systems.
- **Whose intellectual property (IP) is fully owned by the vendor.** The vendor being assessed has developed the distributed scale-out file-based storage solution in-house or obtained the technology through acquisition.

- **That was generally available by September 2021 and generates at least \$30 million in annual revenue.** This is to ensure that the vendor product has at least some level of maturity and market traction.

## ADVICE FOR TECHNOLOGY BUYERS

---

Given that the vendors in this assessment are using widely varying product strategies, an important place to start the evaluation process for an enterprise is to understand which of the different approaches appeal to the enterprise and/or are a better fit for its needs. Do you like the idea of being able to manage block-, file-, and object-based workloads on the same storage system through a unified management interface? Do you prefer unified storage (which can avoid semantic loss issues but will use more storage capacity to provide multiprotocol access to the same data object) or multiprotocol access (which uses less storage capacity but where semantic loss may be an issue)? Are you a federal agency that requires FIPS 140-2 compliant encryption? Do you prefer a storage architecture built around server-based storage nodes or are you open to different architectures that may offer differentiators in certain environments? Six of the vendors assessed use server-based storage nodes (although some of them have some proprietary content), while two – NetApp and Pure – use different architectures.

Would you prefer to use traditional access methods like NFS and SMB but also have access to an intelligent client that offers significantly more parallelization if/when you might need it? Other vendors will tell you how they've extended the performance capabilities of NFS over TCP beyond the 2GBps limit per mount point with nconnect or features specific to their platform that still use the standard NFS client (for example) so you don't have to deploy an intelligent client. Do you require NDMP support? Are you interested in the idea of a cacheless architecture that can offer very high degrees of data concurrency or do more traditional cache-based architectures meet your needs just fine? Do you need POSIX compliance? POSIX really isn't the future, but there are hundreds of thousands of already deployed applications that use it.

Do you have a preference for an HCI-based architecture (like Cohesity or Nutanix) or a disaggregated storage approach? Do you want to buy your solution from a major OEM (Cisco sells Cohesity, Dell sells Nutanix, and HPE sells Qumulo) or would you prefer to buy it from the developing vendor directly (or a channel partner of theirs)? Do you like the idea of combining data protection and enterprise file sharing under a single system or not? While this is not an exhaustive list of questions, these are the kinds of questions an IT manager should ponder when evaluating scale-out file systems for enterprise workloads.

As with most enterprise workloads, high availability (HA) is important and enterprise file sharing is no exception. Solutions that have been around for a long time tend to have an extensive, proven feature set in this area. Understand your recovery point objectives (RPOs) and recovery time objectives (RTOs) for both local and disaster recovery, and match that with capabilities in the scale-out file system offerings. Tunable erasure coding (EC) (so data durability and capacity utilization can be set differently for different workloads), snapshots, replication, a simple "snap to object" feature that makes it very easy to back up the entire namespace to an external object store, air-gap protection to defend against ransomware, and integration with third-party backup products like Commvault and Veritas, all these are features that can impact data protection workflows, availability, and recovery times.

Ease of management at scale is another differentiating area. There are many challenges in managing scale-out file system environments, and there has been a lot of employee interchange between the

various scale-out file system players in the past 20 years. The challenges are well known at all vendors, but how they address them varies. If you have managed a scale-out file system before, what are your hot-button issues?

- Do you need absolutely the lowest latencies for random small file accesses or are sub-millisecond average response times good enough?
- Are you trying to consolidate workloads across your data stage pipelines that need both native and intelligent client-based access methods?
- Do you want to be able to rapidly create delta differentials for backup purposes without having to walk all the file trees?
- Do you want particularly low-capacity utilization of on-disk data protection options at your target level of durability because you have multiple petabytes of data under management?
- Do you need support for compression and/or deduplication because your data sets can benefit significantly from these technologies (or not, since much unstructured data does not compress and/or deduplicate very well)?
- Are disruptive upgrades and slow disruptive recovery in SMB environments a particular pain point?
- Are you particularly concerned about large capacity drive rebuild times or how easy and nondisruptive it is to expand the cluster by adding a new node?
- Are you concerned about how easy and efficient is it to use file quota management systems?

These (and many more) are all issues many scale-out file system administrators have struggled with.

The key to selecting a platform best suited for your requirements is to thoroughly understand your needs and preferences up front. The vendors assessed here all provide a range of performance, scalability, availability, and core functionality that meet the requirements for most enterprise file-based workloads, but among the eight vendors, there are very different ways to get there and very different emphases in their product designs. List what is most important to you, and map that to the vendor offerings. Doing that will require going beyond this document since we do not provide direct head-to-head comparisons between vendors. IDC has, however, published a number of technical reviews of different vendor offerings in separate research, discussing the benefits of the approaches they have taken.

## VENDOR SUMMARY PROFILES

---

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While every vendor is evaluated against each of the criteria outlined in the Appendix, the description here provides a summary of each vendor's strengths and challenges.

### Cohesity

Cohesity is positioned in the Major Players category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

When it was founded in 2013, Cohesity was one of several "new style" data protection vendors that provided a backup appliance that was sufficiently performant and scalable that backup data could be used by a variety of other workloads (analytics, test/dev, etc.) without impacting data protection tasks. Its approach was unique in that it offered an easily scalable platform based around HCI (a relatively

new software-defined storage architecture at the time), and its storage operating environment was architecturally very similar to Google FS (Cohesity's founder Mohit Aron had worked as an early tech leader at Google). Emerging right at the beginning of the digital transformation trend, Cohesity offered a clear way for enterprises to get more value out of their data without requiring a proliferation of storage silos or as much data migration (between platforms). The company grew very rapidly, achieving unicorn status in 2018.

In 2019, the vendor broadened its data management focus with the introduction of a distributed scale-out file system called SmartFiles. While SmartFiles could be deployed as a standalone scale-out file system, Cohesity provided an innovative adoption path for its existing data protection customers. Stored backup data could be exposed as file system "views" available through selected access methods, which then had full read/write capabilities for a scale-out file system running on the vendor's Helios (a trademarked brand name) platform. Helios is based around an HCI architecture, but Cohesity does not compete as an "HCI" vendor – Cohesity bills itself as a data management platform that happens to use an HCI architecture.

SmartFiles is a software-defined storage system that supports significant deployment flexibility in both on-premises and public cloud-based environments. The system uses a file-based data organization method, although it supports multiprotocol access through NFS, SMB, HTTP, and HDFS, as well as access to file-based data through S3 and Swift. Data can be captured from clients through any of these protocols, but the vendor also has a "custom" client (the gRPC) that it uses with its Data Protection Suite product that offers optimized backup and restore operations (and so should be viewed as another data capture method). The gRPC is not an intelligent (parallel) client for file access though.

## **Strengths**

In addition to enterprises attracted to SmartFiles by its capabilities alone, Cohesity data protection customers get the added bonus of a very simple transition to scale-out file system capabilities. The Helios platform (that's what Cohesity calls the core platform it uses across both data protection and scale-out file use cases) presents a potentially significant opportunity to consolidate workloads across use cases onto a single platform, improving administrative productivity and lowering infrastructure costs.

Based on its own merits, SmartFiles differentiates itself by deployment flexibility, comprehensive global data management capabilities, and the ease of use of a "second generation" distributed scale-out file system. With its Google FS "DNA," it supports high scalability on metrics like file size, file system size, number of snapshots, cluster size, and number of defined security roles. Cohesity is one of the vendors that has transitioned to a software-only model, but its customers can get appliances from its channel partners installed on different types of server-based storage hardware. Enterprises can also install SmartFiles in public cloud environments from Amazon, Microsoft, and Google. Its deployment flexibility also allows it to be installed on a virtual machine (VM) for low-cost edge deployments.

A somewhat unique feature of SmartFiles is that its software licenses, all of which are purchased through subscription-based licensing, are portable across different deployment hardware and models (which means, for instance, that if a customer decides to change the underlying server hardware on an upgrade or do a move to the cloud, the customer does not have to relicense the product). While the product has many built-in features that support enhanced performance for random access to millions of small files and sequential access to large files, Cohesity customers highlight SmartFiles' sequential performance.

SmartFiles supports a variety of enterprise-class data services. Administrators can choose between EC and replicas for on-disk protection. Inline compression, inline global deduplication, and software-based encryption (which is FIPS 140-2 certified) can be toggled on and off on a per file system basis. View pinning and quality-of-service controls aid in performance management, and like many other systems, it supports tiered storage within a single namespace, uses AI-driven algorithms for intelligent data placement, and stores all metadata on SSDs for fast access. Built-in data mobility services are also provided to automate data migration from tier 1 NAS devices to SmartFiles and archive data to the public cloud.

SmartFiles supports a very complete set of security features and ransomware detection, prevention, and recovery capabilities. The system supports an unlimited number of space-efficient snapshots (limited only by raw capacity), and snapshots can be defined as immutable, support snapshot "consistency groups," offer a simple "snap to object" feature to external object storage using S3, and can be replicated to remote sites for air-gap protection. The product's ransomware protection is further enhanced through secure data placement (SmartFiles does not overwrite data in place), secure communication via IP white listing, granular RBAC, and very rapid recovery (through integration with its data protection platform). It can also create air-gapped copies to tape.

Cohesity has turned its offering into a "platform," complete with a published API and a marketplace feature. The vendor has developed some additional "applications," available from the Cohesity Marketplace, where customers can download plug-in applications for search, user behavior analytics, data classification, antivirus, ransomware prevention and detection, archive, and other capabilities, and offers a rich set of published APIs so that third parties such as customers or independent software vendors can write apps for its platform. Today, the Cohesity Marketplace includes apps developed by Cohesity as well as third parties, and a percentage of them are free.

## **Challenges**

Cohesity SmartFiles is a solid product, but it has been shipping less than three years. The Helios platform on which it runs, however, has been shipping for seven years. Although it supports NVMe as well as SAS-based storage devices and a variety of different SSDs and HDDs, the product is not architected for high random IOPS in very latency-sensitive environments (although does provide very good sequential read/write performance). Although it does offer standard object-based access methods, customers may face challenges if they want full-featured S3 object storage on the same platform. That said, close to 20% of SmartFiles customers are using both file-based and object-based access methods against their SmartFiles data stores. Customers wanting S3 access to files should ensure that the Cohesity S3 interface as implemented offers the primitives their object-based applications need. Note also that while data written through file-based access methods can be accessed through NFS, SMB, and S3 access methods on this platform, data written via object methods can only be read by S3 on this platform.

Cohesity is primarily known as a data protection vendor, and many enterprises may not be aware of the SmartFiles offering. The vendor has done a good job of adding SmartFiles to its existing Data Protection Suite customer base, but the awareness issue has kept the vendor from bringing as many completely new customers into the Helios fold as it would have liked. Still, there are a number of SmartFiles-first customers, and many of them have gone on to deploy Cohesity's Data Protection Suite as well.

## Consider Cohesity When

With its deployment flexibility, multi-petabyte scalability, global data management, and AI-driven cybersecurity, Cohesity targets batch analytics, media archive, and file-sharing environments where workflows require both file-based and object-based access to the same data. These file-sharing environments include project shares, home directories, machine logs, video surveillance, PACS and medical imaging archiving, historical records and other at least slightly active archives, Splunk SmartStores, and image and/or audio repositories. When Cohesity SmartFiles wins, it's often because the enterprise felt that it provided all the needed unstructured data management capabilities of competitors at a lower cost and/or with more flexible deployment options. The vendor is very clear about targeting workloads in the scale-out file system and object market that require a capacity-optimized solution rather than the extreme high-performance space that requires very low latency.

An enterprise already using Cohesity Data Protection Suite also looking at distributed scale-out file systems should definitely look at SmartFiles.

## Dell

Dell is positioned in the Leaders category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

Dell is a large vendor of IT infrastructure and is one of the huge market players in enterprise storage overall. Isilon was founded in 2001 and introduced the first commercially successful distributed scale-out file system back in 2004. Headquartered in Seattle, the vendor enjoyed a rapid revenue ramp, went public near the end of 2006, and was acquired by EMC in 2010. By the time Dell acquired EMC in 2016, Isilon was already a huge player in the market. In 2021, Isilon was rebranded as PowerScale by Dell and today enjoys a billion-dollar plus annual business.

PowerScale is a multiprotocol storage platform that uses a file-based data organization method. It focuses on using native client access protocols like NFS, SMB, and S3 and has announced no public plans at this time to introduce an intelligent client. The vendor offers all-NVMe, hybrid, and HDD-only nodes, and the nodes can be mixed and matched in a single cluster as needed. Clusters can scale up to 252 nodes to deliver roughly 950GBps of throughput to a single namespace, but the vendor has also scaled the architecture down to provide low-cost, small-edge platforms in both high-availability and non-high availability configurations.

Dell sells PowerScale as a branded appliance under either a capex or an opex (i.e., subscription-based) acquisition model. It is widely used for mission-critical workloads in media and entertainment; electronic design automation (EDA); healthcare/life sciences; financial services; for audio, video, and media serving; across industries for commercial HPC processing; and as a big data analytics platform. The vendor is moving toward a more software-defined model and does offer options where customers can run OneFS, the Isilon/PowerScale storage operating system (OS), on off-the-shelf infrastructure, on premises, or in the cloud. It is available directly from Dell or from literally hundreds of Dell channel partners.

Dell's unstructured data storage portfolio includes three platforms: PowerScale, ObjectScale, and the Streaming Data Platform. PowerScale is the strategic platform for write-intensive native file-based workloads, ObjectScale is for native object-based workloads, and the Streaming Data Platform is optimized for media streaming (and based on PowerScale). The S3 access on PowerScale is intended to make data sharing easier for workflows that access the same data sets across different applications,

but it plans to keep the two (file and object) as separate systems going forward. This is different from the strategies of other vendors like Huawei, Pure Storage, and VAST Data that don't offer object-based platforms that are separate from their file-based ones.

## **Strengths**

The maturity of the PowerScale offering, combined with its extensive enterprise-class storage management capabilities, forms the real strength of the platform. It supports more industry standard protocols than any other vendor (NFS, NFS over RDMA, SMB, HTTP, FTP, Hadoop, NDMP, GPUDirect Storage, IPv4, IPv6, S3, and Swift); offers a true multiprotocol experience where the same data can be accessed through any of these protocols; has more standard features to address enterprise requirements for high availability, fast recovery, ransomware defense, security (including FIPS 140-2 certification), retention (WORM), tiering (CloudPools and SmartPools), and multitenant management capabilities; and includes features that many other vendors do not like compression and deduplication. Customers consider Isilon/PowerScale as a "five-nines plus" platform supporting nondisruptive operations (in almost all cases) and routinely consolidate their mission-critical file-based workloads on the platform. As the original distributed scale-out file system platform and a huge player, it is the standard against which all other vendors in this market measure themselves.

Dell's all-NVMe nodes (the F900s) can deliver very low latencies, but the vendor also offers hybrid and HDD-only nodes that give it the ability to sell to customers that may not want an all-flash offering. Customers often run a variety of different workloads on PowerScale, and OneFS (the storage OS across both Isilon and PowerScale) includes optimizations for environments with millions of small files accessed randomly as well as for environments with large files accessed primarily sequentially. The platform has its own integrated volume manager and uses RAID for on-disk data protection.

Dell's channel partner ecosystem is extremely comprehensive, and the vendor has worked with a number of channel partners to create optimized solutions for many vertical markets. This ecosystem gives the vendor more reach than almost any other competitor (NetApp is probably the closest in terms of size of channel ecosystem). And the vendor's huge installed base, which includes hundreds of thousands of nodes, gives the vendor a power of incumbency on technology refresh that contributes strongly to its market dominance.

Dell Isilon/PowerScale customers often comment that the scale-out file system fits into the vendor's rich ecosystem of management and automation tools, including AI Ops tools like CloudIQ, backup and restore tools like PowerProtect and Veeam (not owned by Dell), data management tools like DataIQ, and data protection OEM solutions like Ransomware Defender and AirGap Enterprise. The platform offers excellent integration with the VMware ecosystem, and many Isilon/PowerScale customers have other Dell storage platforms like PowerStore, PowerMax, and PowerFlex.

## **Challenges**

Customers looking for a true software-defined product that can be deployed across a number of different server-based storage hardware platforms and quickly and easily accommodate new storage device types when they ship will find more software-defined offerings from other vendors. Dell is clearly moving in that direction, and that capability will be a key strategic focus of the new PowerScale line (whereas the older Isilon platform was more hardware defined).

While PowerScale scales out to more nodes than most other vendors (and has validated its scalability to 252 nodes with actual testing), it is not an industry leader in terms of infrastructure density. For

many workloads, competitive systems sized to provide the same capacity and/or performance can be more compact, needing less floorspace, drawing less power, and exhibiting higher storage density (TB per U).

While Dell makes PowerScale easier to manage every year, some of its distributed scale-out file storage competitors can make a strong case that managing their systems requires a far less sophisticated storage administrator. PowerScale's enhanced software capabilities like SmartConnect and SmartPools deliver flexibility and versatility for automated storage and data management that may be lacking in other platforms though. Certain administrative operations, however, can take a bit longer with the platform.

### ***Consider Dell When***

When high availability and the breadth of storage management capabilities in primarily file-based environments are the primary concern, Dell PowerScale is a strong competitor. In addition to covering all the basic requirements across the broadest access method offering, it has additional features that some other competitors don't have but may not necessarily appeal to all buyers. When enterprises prefer to purchase an entire stack solution, including vertical market software and relevant vertical market expertise, PowerScale is available from more channel partners than any other evaluated scale-out file platform. Although PowerScale is not as software defined as some of its competitors, PowerScale does make the platform available in edge, core, and cloud-based versions as well as a managed services offering available from Faction, a Dell channel partner.

With its many features, PowerScale storage pools can be configurable to meet a variety of different I/O profiles, and the system is widely deployed for a variety of different applications: IT workloads like file shares; project and group data; home directories; archive storage and as a backup target; production workloads like real-time streaming; rendering and post-production editing; design automation; DNA/RNA sequencing; genomics; PACS medical imaging and digital pathology; internet/web 2.0 workloads like media serving, audio, video, and web analytics; financial workloads like multiday algorithmic trading and Monte Carlo simulations; and security workloads like camera, drone, and satellite surveillance and feeds. It is also used for big data analytics across a variety of industries, both AI/ML-driven and more conventional commercial HPC-style processing.

### **Huawei**

Huawei is positioned in the Major Players category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

Founded in 1987, Huawei is a China-based vendor of IT infrastructure solutions spanning servers, storage, networking, and telecommunications. Huawei is one of the huge players in the market, and it has a very rapidly growing enterprise storage business that is in the two-billion-dollar range. The vendor has global operations but derives the majority of its revenue from the China market where it dominates pretty much all IT infrastructure segments, except cloud spending. In enterprise storage, Huawei has a broad portfolio that includes block-, file-, and object-based storage platforms and its systems include a very comprehensive set of enterprise-class storage management features. Its broad portfolio includes SAN, unified (block/file/object) storage, and scale-out file systems that can be deployed using either appliances or software-defined solutions and hyperconverged or disaggregated infrastructure. Huawei's solutions can be deployed either on premises or in Huawei's own public cloud environment, and the company offers both capital expenditure and subscription-based purchasing options.

Enterprises considering Huawei storage infrastructure should know that the vendor builds much of its own hardware and at least in the past five years has been an early adopter of industry-changing new storage technologies like NVMe and NVMe-oF. Huawei has two decades of enterprise storage experience and five years ago made a strategic decision to move away from the broad use of commodity off-the-shelf hardware and siloed storage software designs. The vendor firmly believes that the on-premises storage infrastructure of the future will be consolidated in the sense that block-, file-, and object-based workloads will be increasingly run on a single clustered platform that is managed from a unified dashboard.

To that end, the vendor developed a single storage operating system that it uses for its multi-controller SAN, HCI, and scale-out unstructured data storage platforms. This storage OS uses an underlying key value store as its data organization method; supports block, file, object and HDFS interfaces; and offers the configuration flexibility to meet varying I/O requirements in a single system. While some other vendors are also moving in this direction (Nutanix, VMware, etc.), Huawei is farther down this path.

The vendor offers several types of storage hardware that satisfy a range of customer requirements. Two of these systems demonstrate the diversity of the portfolio. The OceanStor Dorado is an all-NVMe multi-controller architecture Huawei primarily sells into the SAN market, while the OceanStor Pacific is its scale-out platform deployed on ARM-based storage nodes. The OceanStor Pacific can be used to support scale-out block-, file-, and/or object-based data stores and can be deployed either as HCI or as a disaggregated storage infrastructure. Both of these systems use a lot of purpose-built hardware. When using it as a scale-out platform, customers will configure it to support the kinds of access and workloads in their environments. Huawei notes that prior experience selling dedicated SAN and scale-out file system platforms in the past helped guide its design of a truly unified storage system.

The vendor's enclosures incorporate unique designs that support more drives per U, augmented cooling that offers more headroom than off-the-shelf designs to accommodate newer, "hotter" storage devices over time (for an extended depreciation life cycle), the vendor's own purpose-built controllers and storage devices (a combination referred to as FlashLink), and a number of optimized software features that transcend traditional data protection and data reduction capabilities. The vendor offers its own version of NVMe-oF (called NoF+) and as a result of all these innovations claims to deliver better performance and efficiencies than industry-standard implementations. Huawei has "open sourced" the NVMe-oF changes, but they are not widely used outside of China.

Huawei's OceanStor Pacific can be deployed as a distributed scale-out file system using NFS, SMB HDFS, and S3 access methods, but it also supports an intelligent "parallel" client (called the DPC) that enables the system's use as a parallel file system. Huawei uses a multiprotocol design where multiple access methods can access a single logical copy of a file or object in the underlying key value store. As mentioned previously, this enables very efficient data sharing and helps keep costs low but customers may risk semantic loss issues. The vendor has introduced some interesting innovations to minimize semantic loss so if this is a concern for customers, they should delve into the technical detail Huawei uses to address the issue. Huawei supports block, file, and object on the same platform, although block data cannot be accessed meaningfully by file and object protocols on the storage system. The capacity for all of these data types is, however, allocated from a single unified pool of storage system capacity.

Huawei's systems support AI-driven monitoring and management. They are somewhat unique in how they size systems for on-premises deployment with customers. While several vendors allow customers

to specify performance, resiliency, and capacity requirements as the way to "buy" managed services, Huawei supports this model for on-premises IT infrastructure purchases. Once customers specify data types, access methods, performance, availability, and capacity requirements, Huawei will suggest the hardware, software, and deployment model to meet that. Objectives are entered into the AI-driven management platform, and the system is managed to meet specified service-level agreements (SLAs) as necessary. As examples, data protection (EC) adapts in the face of failures to maintain defined resiliency levels without having to replace failed devices, and data reduction moves between inline and post-process approaches to maintain defined latencies.

## **Strengths**

Scalability, configuration flexibility, and value for the money are the key differentiators for OceanStor Pacific. OceanStor Pacific claims to support up to 4,096 cluster nodes, but its largest deployed single production cluster has 320 nodes (and 157PB of capacity). The vendor also counts several single clusters of more than 100 nodes in production among its installed base. Innovative metadata and file locking designs support the system's high scalability, and all data services from the vendor's comprehensive list (EC, replicas, compression, deduplication, snapshots, quality of service, encryption, quotas, replication [including "six-nines" rated stretch clusters], etc.) are bundled with the base price of the system and can be configured very granularly. To date, customers largely deploy only a single OceanStor Pacific cluster at a single location, expanding its node count as data grows or they add workloads. While customers can configure a unified namespace in a single system that spans multiple sites, there are customers that have separate clusters in separate sites.

OceanStor Pacific is one of the most versatile offerings in the amount of configuration flexibility it provides (with the exception of its inability to easily deploy the software on non-Huawei servers and non-Huawei public cloud providers). The way data protection can be configured to trade off issues like resiliency versus cost, the option to use mixed media (solid state and spinning disk) and node types in a single system (which supports a very friendly technology refresh model), the ability to support a variety of different access methods to data in the same system, the dynamic ability to shift between inline and post-process data reduction as needed to meet SLAs, and the configuration granularity of the data services (which can be applied at the volume, file, and/or object as well as at the group level), all give customers significant opportunities to meet objectives across all data types in a single system.

Storage systems from Huawei have traditionally been very competitively priced – not necessarily from a list price but from a street price point of view. Bundling of almost all storage management capabilities with the base purchase price of the systems adds to that value. The vendor's early adoption of exciting new storage technologies and the capabilities of its systems (both OceanStor Dorado and OceanStor Pacific) has made it an attractive economic choice for many customers.

## **Challenges**

While the comprehensive unified storage approach used by Huawei has much to commend it, many customers still prefer to keep block, file, and object storage separate. Many systems combine file and object on the same system, and there seems to be a higher comfort level in enterprises with combining those two but keeping block storage separate. If this is a concern, customers should keep in mind that Huawei's storage OS can be deployed as a dedicated block-, file-, and/or object-based platform. It is not an uncommon view among IT practitioners that "jack of all trades" designs are optimized for none, but the reader should note that Huawei's storage OS delivers very high performance across multiple metrics – a single OceanStor Dorado system can handle 21 million IOPS (putting it in the rarefied air of only two other high-end SAN systems, the Dell PowerMax and the Hitachi VSP). Note also that

OceanStor Pacific also turns in very respectable sub-500 microsecond latencies (when using NVMe-based nodes and NoF+ host connections) and can spread single large files across as many as 256 nodes.

While performance, density, and efficiency metrics validate Huawei's choice to use proprietary hardware, the lack of multiple hardware sources may be an issue for some. It is important to note that Huawei is not the only vendor to make this decision though. There are other enterprise storage vendors that claim to provide the same differentiating benefits that Huawei does relative to commodity off-the-shelf storage device and enclosure products to justify their use of proprietary hardware.

The Huawei storage OS supports both self-encrypting drives (that come from Huawei) and software-based encryption. By the feature cutoff date for this evaluation (September 30, 2021), encryption was still only 128-bit, a factor which may have kept it from competing for certain business opportunities, although the vendor does have plans to introduce 256-bit encryption in the future. The vendor does not support FIPS-certified encryption of any type and is not likely to get it. Other areas where the vendor needs to provide more capabilities to be on par with most of its competitors are in tiering to external targets and in providing improved multipublic cloud support.

Huawei has hundreds of millions of dollars of OceanStor Pacific revenue and is in the Major Players category competing in the scale-out file system market. Still, the fact that most of its revenue comes from the China market may be a concern for some prospects. The vendor does have a worldwide 24 x 7 support capability, but only a very small percentage of its revenue comes from Europe and the Middle East, and none of its revenue comes from the United States (which is the world's largest enterprise storage market).

### *Consider Huawei When*

Enterprises that buy into the comprehensively unified storage approach with block-, file-, and object-based storage all running on a single, highly scalable storage platform will definitely want to look at Huawei OceanStor Pacific. There are only two vendors that support both native access methods and a true parallel client, and Huawei is one of them. There are also only two vendors in this assessment that support block-, file-, and object-based storage on the same platform, and Huawei is one of them. This provides more opportunities for workload consolidation across a wider range of I/O profiles, driving a potentially advanced economic value proposition. In environments with multistage data pipeline workflows, the ability to run high-speed ingest workloads using the DPC on the same platform where NFS and/or SMB may be used by workloads in later pipeline stages will be able to more effectively share data (and possibly reduce the number of separate storage platforms that must be deployed to handle an enterprise's big data analytics requirements).

Enterprises not as concerned about proprietary hardware platforms and more interested in the better performance and scalability, increased efficiency of resource utilization, and improved media endurance of the vendor's all-flash storage should also take a look at Huawei OceanStor Pacific.

### **NetApp**

NetApp is positioned in the Leaders category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

NetApp is a large enterprise storage vendor founded in 1992. Up until NetApp appeared, most enterprises were using standalone Unix, Windows, or Linux servers for file sharing. As file-sharing needs grew, file servers proliferated and very quickly became extremely difficult to manage. NetApp

created a family of "filers" based around a dual-controller, enterprise-class design that were optimized for file serving, could be quite large, and were highly available. Many customers could consolidate their tens or hundreds of file servers onto one or just a few NetApp filers. This was hands down a better approach to enterprise file serving, and the vendor enjoyed a very rapid rise throughout the 1990s.

Enterprise file workloads began changing, however, in the early 2000s, and enterprises found their data sets growing significantly. Distributed scale-out file systems that could spread a unified namespace across many nodes in a cluster appeared, offering a scalability path that was easy to take beyond the capabilities of even the largest NetApp filer. As scale-out architectures began to take hold over the decade, NetApp introduced an ability to cluster its dual-controller arrays in a single namespace, providing the ability to scale a file system namespace across cluster nodes (but no single file could span multiple nodes). IT managers should note, however, that "clustered scale up" (which is what NetApp does) is architecturally different from "distributed scale out" (which is what Isilon introduced and many of today's competitors use).

This short preface is important because NetApp offers customers many advantages with its approach while not being able to do other things that distributed scale-out file storage platforms can do. How a configuration is scaled, and the performance implications of that, differentiates the two types of offerings, and customers should understand this. Even though it offers a unified namespace across multiple nodes in a clustered scale-up system, a single file cannot very easily be distributed across multiple nodes, which means that the access performance to that file is determined by the single node that it is sitting on. In a distributed scale-out file system, a single file is generally automatically spread across multiple nodes, distributing the workload more evenly and potentially smoothing out I/O spikes that may occur. In practice, however, there may not be much difference in performance for some workloads, depending on file size and access profiles, whereas for others there may be a noticeable difference.

NetApp's Data Fabric, the vendor's hybrid multicloud strategy, is probably the most complete in the industry. ONTAP (the vendor's enterprise storage operating system) can be deployed as appliances on premises and as software-only in public clouds and can be purchased as a first-party service offering from Amazon, Microsoft, and Google that actually runs on NetApp hardware to deliver better performance, availability, and functionality than native file services offerings from the hyperscalers. NetApp's Cloud Manager provides monitoring, administration, and orchestration for ONTAP storage on premises and in the cloud, built upon a comprehensive and consistent set of enterprise-class management capabilities across all deployment models. And Cloud Manager is well integrated with the vendor's AI Ops platform, Active IQ, which uses AI/ML-driven analytics to help improve systems management, performance, availability, and capacity planning operations.

NetApp's filers run the ONTAP storage OS and today support unified (block- and file-based) storage. ONTAP uses a hierarchical data organization method (referred to as WAFL) but can store both block-based and file-based data in it. NetApp systems are deployed in HA pairs, and a single federated cluster can accommodate up to 24 nodes of different sizes and generations (all of which must be running ONTAP). NetApp offers a good variety of different system types, with all-NVMe, hybrid, and HDD-only arrays. The all-NVMe systems are branded as All-Flash FAS (AFF), and the hybrid and HDD-only systems are branded as FAS (and use 12Gb SAS). They are available in a variety of different sizes, each with different performance and capacity points, but all run the same storage OS.

NetApp not only offers a variety of different access methods on its ONTAP arrays, including S3, but also offers a separate object-based storage platform called StorageGRID (obtained through the

acquisition of Bycast in 2010). The vendor positions StorageGRID for native object workloads, while the ONTAP systems enable data sharing between file- and object-based workloads. The vendor also offers a parallel scale-out file system based on BeeGFS but uses its E-Series arrays for the hardware (the E-Series runs the SANtricity rather than the ONTAP storage OS). Most of the competitors evaluated do not offer a separate parallel file system platform.

## **Strengths**

NetApp has almost the perfect architecture for more traditional enterprise file-sharing workloads. The fact that it is based on dual-controller arrays that come in a variety of sizes makes it very easy to deploy; each system is very highly available and offers extensive enterprise-class storage management capabilities that are proven. The system offers features somewhat unique in the scale-out file system market like stretch clusters that have a zero RPO, and it is widely used for mission-critical workloads. It has built-in data protection, a comprehensive set of snapshot capabilities, and optional ransomware prevention, detection, and recovery features. It is easier to deploy a single HA pair than a multinode cluster (which is primarily what the other distributed scale-out platforms use), and customers can choose from small to very large systems so they can deploy the platform most accurately sized to their requirements. ONTAP offers many features that support six-nines plus availability, fast recovery, nondisruptive upgrades, and great ransomware protection.

The ONTAP arrays support all-NVMe configurations as well as NVMe-oF host connections (Fibre Channel, RoCE v2, and TCP). For most workloads, the performance of the systems means that "hot spots" are generally not an issue for single files (which can be up to 16TB in size). NetApp's high-end array, the AFF A900, delivers over 1.4M IOPS, 25GBps of bandwidth, and 2PB of effective capacity (assuming a 3:1 data reduction ratio) in just 4U. (Note that the vendor offers inline compression, data deduplication, and encryption capabilities. It supports both hardware- and software-based encryption options, as well as cryptographic modules that have been FIPS 140-2 validated.) A fully configured 24 node cluster can support over 700PB of effective capacity (which would be spread across multiple namespaces). Storage clusters based on server-based storage nodes would need many more nodes to deliver this same performance. ONTAP supports 400 billion files and up to 20PB in a single namespace.

Finally, the vendor has a huge installed base of ONTAP systems that give it the power of incumbency when it comes time for technology refresh. As a vendor, NetApp has generally introduced newer technologies faster than many (although not all) of its competitors. The new executive team brought on by new CEO George Kurian in 2015 has made many positive changes to the original NetApp's strategy, one of which was to target generating more of its quarterly revenue from new rather than existing customers. Since 2019, NetApp has made major strides in broadening its appeal to new customers (upping the quarterly rate of new customer acquisition) and improving the customer experience (CX) it offers customers. The Simplicity365 program has comprehensively impacted a number of customer touch points to make everything from generating a sales quote and making deployment and provisioning easier to improving operational efficiency and speeding problem resolution in technical support situations. The vendor tracks NPS internally and has been enjoying improving scores, but it does not publish this data externally.

## **Challenges**

While NetApp does sell its storage OS as a software-only product for on-premises or public cloud-based deployments, the vendor does prefer on-premises deployments to use its appliances. The NetApp hardware is not the same as commodity x86 server-based storage hardware, and while it

provides more effective redundancy and performance, it is more expensive. Depending on which NetApp nodes a customer starts and expands with, it will need fewer nodes (potentially far fewer) than competitors using server-based storage architectures, but the nodes will be more expensive.

The original NetApp architecture was less well suited for the bigger data sets in file-sharing workloads that emerged in the early 2000s, but the introduction of its scale-up clustering capabilities did improve its ability to handle these environments. Still, there are workloads and workflows where the inability to easily stripe data in a single file across many nodes can be a limitation, and the architecture cannot deliver high data concurrency in the same way that some scale-out file systems can (which can be a potential limitation with some more modern application workloads). Also, deduplication is performed at the aggregate level within a single node (an aggregate is a collection of one or more RAID groups) so customers do not have the benefit of global deduplication. However, NetApp does have a post-process deduplication capability called Temperature Sensitive Storage Efficiency that identifies cold blocks from which it wrests additional space savings in each node. Because of these design considerations, there are some enterprise workloads where other architectures can offer potentially better performance, scalability, and ease-of-use benefits.

While NetApp has definitely improved the usability of ONTAP over the decades, it assumes a level of storage management sophistication above some of its competitors. This approach uses fewer defaults but gives IT operations personnel more opportunity to tune systems to meet very specific requirements. Some enterprises want this level of configurability, while others feel they do not need it and like the ease of use of systems that use default configuration assumptions more comprehensively.

### ***Consider NetApp When***

For those requiring advanced hybrid multicloud capabilities, NetApp has an extremely comprehensive set of features and an excellent vision. NetApp was the first major vendor to realize the importance of hybrid cloud support and has been enhancing its capabilities in this area for longer than other vendors. NetApp's customers firmly believe the vendor leads the industry with its hybrid multicloud support.

ONTAP-based systems offer very good random read/write performance against small files and can offer good performance against larger files as long as a single ONTAP array has sufficient bandwidth to handle all accesses (since a file cannot be spread across multiple HA pairs very easily). NetApp wins when it comes to supporting high-availability file-sharing environments and has a very strong channel ecosystem with much vertical market expertise. When most file workloads for an enterprise can fit on one or just a few HA pairs supporting a single namespace, NetApp is very easy to deploy, but as performance requirements that exceed the capabilities of a single node arise, some of the inherent capabilities of distributed scale-out file systems can become more important. These scale-out limitations are not as much of an issue for workloads that work just fine with multiple namespaces in a single node but having multiple namespaces does make it more difficult to search, analyze, and protect the data. This is why some competitors tout their very large file system sizes.

### **Nutanix**

Nutanix is positioned in the Major Players category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

Founded in 2009, Nutanix is a publicly held large HCI vendor targeted for enterprise use. As one of the vendors that established the HCI market, it distinguished itself with customers both for its deployment flexibility and ease of use. It was the original proponent of "one-click operations," a significant

improvement for many common IT workflows. While many of its competitors have introduced some one-click operations as well, Nutanix has maintained its differentiation in both of these areas. Starting as a block-based vendor of appliances, Nutanix initially required the VMware ESXi hypervisor but then added its own KVM-based AHV hypervisor option (marketed as helping its customers avoid the "VMware tax") as well as other hypervisor options. It also added bare metal and more software-defined deployment options across a variety of server-based storage hardware (including an ability to deploy Nutanix either in an HCI or a dedicated mode to deliver high capacity for large-scale deployments). With its ability to support a variety of different workload types, Nutanix provides a very easy to scale and manage platform for enterprise workload consolidation.

In 2017, the vendor added support for Nutanix Files, allowing workloads that used NFS and SMB access methods to also be consolidated onto the platform. Customers add file services by enabling the Files service, which deploys File Service Virtual Machines (FSVMs) to support both NFS and SMB access methods. In 2018, Nutanix moved to a software-defined model, still delivering appliances through its OEM and channel partners. In 2019, the vendor introduced Nutanix Objects, creating a unified (block, file, and object) storage platform that offered even more opportunity for workload consolidation with centralized management across all of an enterprise's workload types. The vendor's tagline is "simple, secure, and scale out" and while that is accurate, the vendor's solution (now referred to as the Nutanix Cloud Platform [NCP]) is also notable for its compelling cost-effectiveness compared with more traditional three-tier architectures.

Its software-defined architecture makes Nutanix a natural fit for cloud-based deployment, and the vendor offers its customers a good hybrid multicloud strategy. Pricing for its software is subscription based, and customers can obtain it either for on-premises deployment or for deployment in Amazon. Cold data in Nutanix Files can be tiered to any S3 compliant external storage (including that running on any public cloud provider or Nutanix Objects). Features like Nutanix Cloud Clusters (NC2) make it easy to move workloads in VMs to the public cloud (or back). Nutanix Files can also be deployed in AWS (via NC2) to provide customers the flexibility to use it in a hybrid cloud model.

## **Strengths**

The vendor's unified storage strategy offers very attractive capabilities as well as an easy adoption path for customers interested in consolidating multiple access methods onto a single storage platform. It is very easy for existing Nutanix HCI customers to add file- and/or object-based capabilities to their existing deployments in a nondisruptive manner. The vendor has been very successful selling Files to existing customers – since the product's initial introduction in 2017, the vendor has consistently grown its customer count by 60% year over year. While that is an easier sell for it as a vendor, it also drives a good rate of new customer acquisition. Over the past several years, new customer logos have typically represented in excess of 40% of Nutanix Files revenue each quarter.

Deployment flexibility is a key differentiator for Nutanix. Customers can choose from all-NVMe, hybrid, or all-HDD nodes as well as from extremely storage-dense nodes (which are especially relevant for very high-capacity environments). The vendor supports running Files in two modes, either alongside existing VMs in a consolidated mode or on a separate dedicated cluster so that administrators can scale FSVM resources independently. In either model, Files can run alongside other storage services including both volumes and objects to provide consolidated storage clusters. As with other Nutanix offerings, Files is managed through Prism Central (which is bundled with the base purchase). Prism Central provides unified management of all Nutanix services, storage, or otherwise. The vendor supports snapshot-based data protection that offers user-driven self-restores with file-level granularity

as well as share-level replication through SmartDR for disaster recovery. Immutable snapshots for ransomware protection are also supported.

The ability of Nutanix to start small and grow to multi-petabyte capacities has been a significant advantage, particularly against other storage systems that force larger starting points in terms of capacity. Nutanix Files leverages NCP capabilities to seamlessly scale infrastructure both up and out and provides the ability to scale performance and capacity independently (through the dedicated mode mentioned previously) – a feature which allows it to accommodate higher-performance workloads than more traditional HCI architectures. Nutanix can start as small as a single node and grow to a namespace that can span as many as 48 physical nodes. Single- and dual-node configurations enable lower-cost deployments for test/dev as well as ROBO use cases, while highly available enterprise systems require three or more FSVMs, each of which must be running on a different physical node.

The vendor takes security very seriously. In addition to supporting FIPS 140-2-certified encryption for data both at rest and in flight (using NCP's replication feature), the vendor offers role-based access control and integrated security analytics to prevent, detect, analyze, and recover from ransomware and other security threats. Integrated ransomware protection uses a growing database of over 4,000 ransomware signatures, while an anomaly detection engine monitors data for potential threats such as mass file deletions, mass permission changes, mass encryption efforts, and other potentially suspicious activities. Nutanix Data Lens is a cloud-based data governance service that delivers actionable insights into access patterns, data age, types, and other contextual information to enable efficient data life-cycle management, improved protection against ransomware attacks and insider threats, and "proof of compliance" for regulatory and/or governance purposes.

Last, the vendor has distinguished itself as the enterprise storage vendor that has the highest published NPS. Even as it has grown to have over 20,000 customers overall (across its entire product line), it has maintained an NPS of 90+ over the past seven years. This is unheard of not only in the enterprise storage industry but across almost all of the 220+ industries that use NPS as a measure of customer experience. The vendor notes that this performance has a significant influence on new opportunities, expansions, and renewals. This is without a doubt the highest published NPS in the enterprise storage systems market.

## **Challenges**

While Nutanix has continued to grow revenue over time, profitability has been elusive. Rajiv Ramaswami came on as the new CEO in late 2020, replacing Dheeraj Pandey who had announced his pending retirement earlier that year. In early 2022, Ramaswami announced that Nutanix was on a path to profitability by the end of calendar 2022. There doesn't appear to be any near-term financial risk in deploying Nutanix, and if the vendor hits its profitability goals, this may completely remove this as a concern.

As the vendor introduced so many new capabilities in its platform over time, the product portfolio grew, and licensing became more complex. While Nutanix customers are very happy with the Nutanix Cloud Platform (note the current NPS of 92), there was a need to make licensing easier and address co-termination issues. In early 2022, Nutanix implemented a significant simplification in these areas that promises to make the Nutanix Cloud Platform as easy to buy as it is to deploy and use. If executed well, this program should help boost revenue growth.

While HCI is the fastest-growing segment of the software-defined storage market, it has achieved its greatest success in small and medium-sized enterprise. While HCI boasts strong flexibility, simple

ease of use and scalability, and compelling economics, there are a few HCI clusters with more than 15 or 20 nodes for most of the vendors playing in this space. Today, Nutanix does have customers with 2-4PB of storage deployed in production environments. New multicore processors as well as denser memory and storage continue to increase the capabilities at the node level, but IDC has noted that enterprises that require larger systems tend to move to disaggregated storage architectures that give them more leeway in scaling compute and storage resources independently. While this does not impede Nutanix from selling to large enterprises, it does limit the vendor in some cases from selling to customers that want to deploy larger storage systems.

### **Consider Nutanix When**

Nutanix Files stands out for its simplicity and ease of management and expansion, and for customers that want to deploy a system that will handle petabytes in single digits across its life cycle, it has much to commend it. Enterprises that expect to support 10PB+ either at initial deployment or across a scale-out platform life cycle may want to ensure that Nutanix Files capabilities meet their requirements in terms of maximum file and file system sizes and other metrics like snapshots per share.

Major customer use cases for Nutanix Files include enterprise file sharing for end-user computing (including VDI), storage consolidation for both legacy and modern applications (including container-based workloads), audio/video repositories, and medical imaging. Nutanix Files is widely deployed in enterprises of all sizes across government, financial services, professional and technical services, manufacturing, healthcare and life sciences, technology and telecom, retail/wholesale, and energy vertical markets. The vendor sells across all geographic regions with both 12 x 5 and 24 x 7 support with next day onsite response capabilities; sells through major storage OEMs like Dell, HPE, Fujitsu, Inspur, and Lenovo; and offers a growing set of reference architectures covering solutions such as end-user computing, video surveillance, and AI. Some of the storage OEM vendors offer pay-per-use consumption models for Nutanix technology (e.g., HPE GreenLake and Lenovo TruScale).

While Nutanix Files offers compelling simplicity of deployment and ease of use, it is primarily a file-based offering. Enterprises that want to enable S3 access to file-based data without having to copy it to a separate storage pool may want to look elsewhere. Nutanix Objects does run on the same NCP as Nutanix Files but is serviced out of separate storage pools from file-based services.

### **Pure Storage**

Pure Storage is positioned in the Leaders category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

Founded in 2009, Pure Storage is a large, publicly held enterprise storage vendor that sells only all-flash storage. The vendor changed the industry with its block-based FlashArray (which originally shipped back in 2012), and in 2016, it entered the unstructured storage market with its FlashBlade (which supports both file-based and object-based storage in the same system). FlashBlade has been very successful for Pure Storage, and had it been an independent business, it would have achieved unicorn status several years ago and actually crossed the \$1 billion in lifetime sales line in June 2021. Even as the industry experienced a downturn during the pandemic years, Pure Storage was able to turn in steady revenue growth and has an installed base of 10,000 customers (across its entire enterprise storage portfolio).

FlashBlade is a unified (rather than a multiprotocol) storage platform supporting NFS, SMB, and S3 access methods and using an underlying key value store data organization method. It is fully hardware

redundant, offers hot-plug replacement of all components, and delivers very high availability in production usage. It supports a broad range of data services, including compression, global EC (which can be spread across all blades across chassis in a cluster), replication, snapshots (including immutable SafeMode snapshots), audit logs, and 256-bit encryption. The system's efficient all-flash design requires less power and cooling and offers higher infrastructure density (for both performance and capacity) than many competitors. This allows smaller, more compact systems to meet customer requirements, also saving on datacenter floorspace. FlashBlade can be monitored and managed by Pure1 META, the vendor's AIOps hybrid cloud management platform.

## **Strengths**

One of Pure Storage's original design tenets, reflected in both its FlashArray and its FlashBlade platforms, is that an all-flash array (AFA) overcomes so many of the performance issues in hybrid and HDD-based storage systems that defaults can be widely used when deploying systems (dual-parity RAID, always-on compression and encryption, etc.). This makes the FlashBlade extremely easy to use, a feature consistently noted by its end users. There are a very few settings which administrators can configure, but there are a very few times when an administrator may want those capabilities (unlike some other systems which often require sophisticated manual tuning expertise). FlashBlade's ease of use extends from initial deployment and storage provisioning to system expansion, upgrades, and failed component replacements.

The other strength of FlashBlade is its ability to deliver consistent performance at scale. FlashBlade uses a cacheless, scale-out architecture that requires significantly less external cabling than clusters built from server-based storage nodes. FlashBlade is a pluggable blade-based scale-out architecture. Its storage devices are called "blades," available in 17TB and 52TB capacities, and each chassis can accommodate up to 15 blades. The blades are a proprietary design – not off-the-shelf SSDs – and include both performance and capacity resources. The flash media on each blade is managed globally by Pure's Purity//FB storage operating system, and these devices have very different and better performance, endurance, and overprovisioning profiles than off-the-shelf SSDs. The internal backplane is based on Ethernet, all the blades are directly connected using the NVMe protocol, and as a result, the system delivers a higher degree of concurrency (an issue particularly important in dealing with large data analytics and densely consolidated storage workloads) than most competitors.

Contributing to FlashBlade's performance characteristics is its scale-out metadata architecture (based on a variable block metadata engine and distributed transaction database), a design which enables it to handle billions of files and objects with equally good performance for small and large files as well as random and sequential access.

CX has always been a focus for Pure Storage, and it has published the industry's only independently validated NPS for five years now. It is notable that its NPS has consistently been in the mid-80s over this period, in particular because for most enterprise storage vendors the quality of their CX tends to degrade as their company grows. Many factors contribute to Pure Storage's CX across the entire product life cycle, including its online sales quoting system, its self-service management interface, the infrastructure density of its arrays, the ease of use in managing its all-flash systems, the consistently high quality of its technical support, and its Evergreen Storage program.

## **Challenges**

FlashBlade uses proprietary hardware. While the vendor makes good arguments about the benefits it offers to customers, that may be an issue for some enterprises. And while the performance it delivers

across its access methods is very good, it does have a limited set of them. Many other distributed scale-out file storage platforms offer a broader array of access method options. Pure Storage is committed to native access methods and has no plans to introduce an intelligent client. While the FlashBlade architecture can handle high data ingest rates using NFS, SMB, and S3, it cannot compete with the "throughput to a single large file" performance of parallel scale-out file systems.

While FlashBlade can tier to external HDD-based storage using S3, it does not directly support HDDs. Although Pure Storage has features that lower the cost per gigabyte at the system level, it is not one of the less expensive distributed scale-out storage platforms to buy, but when it comes to all-flash systems, IDC strongly suggests that it is most important to look at overall total cost of ownership rather than just initial purchase price. Its cost profile makes it less suitable for colder storage workloads, although the advantages of its all-flash design are evident for workloads that have any kind of performance sensitivity, whether that is in supporting high degrees of data concurrency or very rapidly moving large data sets. Roughly 25% of FlashBlade's customers use it as a backup repository, citing its write ingest, infrastructure efficiency, and rapid restore advantages.

FlashBlade is not quite as "cloud friendly" as FlashArray, the vendor's dual-controller array. While FlashBlade supports file- and object-based replication to cloud-based targets, Purity//FB is not available in a software-defined version that can be run in the public cloud. And the platform lacks a deduplication feature, although deduplication does not provide much value for many unstructured data workloads. FlashBlade also does not support NVIDIA's GPUDirect Storage API, although it does offer a converged infrastructure stack offering (AI-Ready Infrastructure [AIRI]) with NVIDIA that includes NVIDIA DGX accelerated compute servers, FlashBlade storage, and Mellanox NVMe-oF networking – all under a single purchase SKU and with a single point of support contact with Pure Storage. The vendor's strategy with FlashBlade is to use industry-standard protocols as much as possible and argues that its high all-flash performance and high degree of data concurrency allow it to do an excellent job of keeping GPUs fed with data (without requiring the use of proprietary interfaces like GPUDirect Storage).

### ***Consider Pure Storage When***

FlashBlade excels at delivering high-performance, high infrastructure density and ease of use for unstructured data storage environments. Enterprises with FlashBlades also comment on the system's ability to densely consolidate workloads with differing I/O profiles, a capability enabled by the high data concurrency it supports. Simultaneous use of FlashBlade as both a backup appliance with rapid restore and a platform for big data analytics projects is very common in its installed base. Top verticals generating revenue for FlashBlade include financial services, professional, technical and business services, government (FlashBlade is FIPS 140-2 compliant and has a thriving federal business), healthcare and life sciences, research, electronic design automation, and advertising, media, and entertainment.

### **Qumulo**

Qumulo is positioned in the Leaders category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

Founded in 2012, Seattle-based Qumulo is a distributed scale-out file system vendor that has been shipping its Qumulo Core file system platform since 2015. Although privately held, Qumulo reached unicorn status with a \$1.2 billion valuation in July 2020 and its recent FY 4Q22 results indicated 75% sequential growth over its third quarter results. The vendor completed its sixth round of funding in July

2020, bringing the total raised to \$351 million. Qumulo sells its file system software on a subscription basis. Nearing 1,000 customers, the vendor is significantly benefiting from digital transformation in enterprises struggling to implement a more efficient strategy for unstructured data management at scale.

## **Strengths**

Differentiators that define Qumulo include flexibility of deployment (because of its software-defined design), simplicity and efficiency of management at petabyte scale (because of features specifically designed to address well-known scale-out file system issues with first-generation distributed scale-out file system platforms), and real-time analytics (that provide comprehensive visibility into file system metrics that enable more effective management of large-scale environments).

Qumulo Core is a software product that can be deployed on commodity server-based storage from a variety of different suppliers and is available as an appliance from channel providers like Arrow, Fujitsu, and HPE and as a cloud-based file storage service from AWS, Microsoft Azure, and Google Cloud. Its software-defined nature supports mixed storage cluster configurations with NVMe and hybrid nodes, delivers exactly the same functionality in on-premises or public cloud-based deployments, and easily accommodates new storage device types and multigenerational technology refresh. Qumulo cloud instances (referred to as Cloud Q, while its on-premises deployments are referred to as Server Q) differentiate themselves from native cloud-based file services on scalability, multiprotocol support, and data visibility (through their onboard real-time analytics capabilities). The vendor offers a unified management console that spans on-premises and off-premises deployment locations for easier management in hybrid multicloud environments.

Architecturally, the system's scalability in terms of file sizes and file system sizes is impressive. Although the company has not validated it for use in production, it claims that its architecture supports single file sizes up to 9EB and unlimited file system sizes. Qumulo does certification work for production use based on customer requirements though, and to date, the vendor has certified a 40PB cluster size (all in a single namespace). These numbers stand in stark contrast to file and file system size limitations of some of the vendors assessed in this document.

While all scale-out file system vendors claim ease of management, Qumulo can point to a variety of specific features that are very attractive to administrators already well versed in managing these types of environments – how quota management is handled with an in-band method that is much more efficient at keeping it in sync with the actual file system, how namespacewide delta differentials are generated within 15-30 seconds for backup purposes, the use of heat metric-based intelligent data placement to optimize cache hit rates, and a number of optimizations that improve the system's ability to efficiently handle billions of small files. (Note that it is too detailed to list here but definitely an area that should be explored by technical decision makers for whom that is important.)

On-disk data protection is purely software defined and operates at the block (rather than the file) level using a flexible EC-like approach (the vendor's platform includes an integrated volume manager/file system layer). This not only enables much faster drive rebuilds but lets customers mix and match different device sizes in the system as well. In addition to these capabilities, Qumulo supports a variety of other features like snapshots, replication, changed block tracking that does not require B-tree walks, air-gap protection to defend against ransomware, and integration with third-party backup products to support high availability and fast recovery. Qumulo is widely used in its installed base for mission-critical file-based workloads.

Qumulo is fully committed to using native access methods like NFS, SMB, and HTTP instead of intelligent clients because of the ease-of-use differences. It uses a multiprotocol access design that maximizes capacity utilization, uses cross-protocol permissions and identity management, and uses 256-bit encryption for all data at rest. Replicated data in flight is encrypted, as is all access using the SMB protocol. The vendor is currently in the process of obtaining FIPS 140-2 certification from the U.S. government. Rich REST APIs enable easy workflow automation integration.

## **Challenges**

While no longer a start-up, Qumulo is still a relatively small player. All Qumulo fulfillment flows through channels, but the vendor does provide a single point of support contact for all its different appliance-based models. Alternatively, customers can also buy its product through established storage OEMs like HPE if there is concern about the vendor's size.

Some features that other vendors offer are missing. Qumulo does not yet support compression and data deduplication, and while it is likely to introduce compression at some point, Qumulo does not intend to introduce dedupe – primarily because the data sets its customers manage rarely benefit from it. S3 is not a supported access method, although the vendor plans to remedy that by the end of 2022. Some multitenant management features like access zones (which allow administrators to partition a cluster into multiple virtual containers to isolate data) are not supported, so customers may want to closely review the feature set to ensure that Qumulo's capabilities meet their requirements.

Qumulo is a firm believer in software-defined systems and tends not to support hardware assist for features like encryption (although it does support that in its HPE Apollo-based appliance). And while it is not necessarily a challenge, Qumulo does not aim to offer the lowest latencies, the highest throughput to a single file or a platform well suited for cheap and deep cold archives – it targets the large middle of the file storage market with a system that offers very flexible deployment options, is easy to manage at scale, and provides comprehensive data visibility to inform better data management.

## **Consider Qumulo When**

Qumulo Core is a good fit for not only more traditional file sharing but also the newer big data analytics workloads being deployed as part of digital transformation that use NFS and/or SMB. Enterprises with hybrid multicloud environments will appreciate the deployment flexibility and cloud support capabilities, and administrators with prior experience with distributed scale-out file systems will appreciate many of the ease-of-use features (including in particular the visibility enabled by the Qumulo Aware component) that really prove their worth as data under management scales to a petabyte and beyond. Key verticals where the vendor has enjoyed significant success include advertising, media and entertainment, manufacturing, technology, software and telecom, health and biotechnology, and professional, technical, and business services. Qumulo already also has a good federal business but expects that to ramp further once it achieves FIPS 140-2 certification by the end of 2022.

## **VAST Data**

VAST Data is positioned in the Major Players category in the 2022 IDC MarketScape for worldwide distributed scale-out file system.

VAST Data is a privately held storage software vendor headquartered in New York City. It has a unique "disaggregated and shared everything" (DASE) architecture that differentiates it from all of its competitors (most of whom are using more traditional shared-nothing architectures). VAST Data is an

all-flash system, combining a thin layer of storage-class memory with high-capacity QLC-based SSDs in a single tier. The SCM tier addresses the write performance and endurance concerns with QLC and drives the performance characteristics of the system, while the QLC (which comprises 97% of the capacity of a system) dominates the cost per gigabyte. The vendor uses single-ported QLC devices for lower cost but provides a high-availability (five-nines plus) system based on other aspects of its architecture. Storage controller enclosures (called C Boxes) are stateless and separate from storage capacity enclosures (called D Boxes), allowing each resource to be scaled in a completely independent manner as needed.

All state is kept in the SCM devices which are in the capacity enclosures, a fact that offers major benefits in reducing east-west traffic in the system. Compared with shared-nothing architectures, there is no lock management, cache coherency, or rebuild traffic that has to go across a "cluster interconnect," enabling the VAST Data Universal Storage Platform to scale significantly better than many more traditional shared-nothing architectures. VAST Data's scalability claims are no idle boast – while the average cluster size among other vendors is in the several petabyte range (it is smaller for some), VAST Data's is closer to 10PB. It is true that other vendors have examples of larger clusters, but it is also true that back-end network congestion limits the size of competitive distributed scale-out systems to generally less than 10PB in production environments (there are a few outliers from that limit for most vendors though). VAST Data has tens of customers in its installed base that are well beyond this with roughly only 200 customers. VAST Data's architecture requires some study before its benefits become apparent – those interested in taking a deep dive into the VAST Data Universal Storage Platform, see *VAST Data: A Technical Deep-Dive Look at a Compelling New Scale-Out Storage Architecture* (IDC #US48805222, February 2022).

VAST Data customers primarily use the system with NFS, SMB, and/or S3 access methods. Its data organization method is hierarchical but is not a standard file system design. VAST Data gives its customers the option to write and read data natively using the same protocol or provide multiprotocol access to data. This is an option where customers can decide on a data set basis whether capacity utilization or semantic loss is the bigger concern. While VAST Data offers options to improve NFS performance with NFS over RDMA or using nconnect, it also will be offering an intelligent client option for customers later in 2022.

Since shipping its initial product for revenue in early 2019, the vendor has turned in extremely impressive financial performance. It experienced over 350% revenue growth in 2020 and became a unicorn with a \$1.2 billion valuation when it completed a C round that brought total funding to \$140 million in April 2020. It ended that fiscal year with over 100 customers. VAST Data achieved unicorn status faster than any other storage vendor ever, was cash flow positive by the end of 2020, and grew revenue in 2021 by 380%. A D funding round completed in May 2021 brought in another \$80 million, driving the vendor's valuation to \$3.7 billion. VAST Data sells into a number of verticals, including financial services, health and biotechnology, government, advertising, media and entertainment, manufacturing, and technology software and telecom, and counts three of the hyperscalers among its customer base.

## **Strengths**

Infrastructure efficiency and price/performance are the aspects of VAST Data's platform that impress customers the most. Smaller VAST Data systems are generally required to meet a given set of performance and/or capacity requirements, and very innovative on-disk data protection (based on a modified form of EC) and data reduction make extremely effective use of raw capacity. This means

that these systems require less energy and floorspace. Performance scalability is also very good, primarily because the addition of front-end clients does not add to back-end network congestion like it does with shared-nothing architectures (and it is well known that back-end network congestion can be a major issue with these more traditional designs).

The vendor's disaggregated architecture offers a number of benefits. First, since any C Box can talk to any D Box, the system stays up and running, with access to all data, even if there is only one C Box left. The vendor supports up to 10,000 C Boxes in a single system. D Boxes actually contain four storage capacity nodes, all of which are connected over a PCIe switch embedded into the enclosure. Each node includes a small amount of SCM capacity with the rest made up of NVMe-based QLC SSDs. For readers potentially concerned about a write cliff, it is important to note that data is written to the SCM in each D Box across two 100GbE links, but destaged to QLC SSDs across four PCIe Gen 4 links (each of which offers up to 512Gbps of downstream bandwidth), a design which basically means that the PCIe switch can't ingest enough data to cause a write cliff in a D Box.

VAST Data originally used 15TB QLC SSDs but later introduced 30TB QLC SSDs, so the system can support significant storage density (TB per U). When a drive fails, rebuilds are very fast and rebuilds get faster as more nodes are added. All state data is kept in the persistent SCM in the D Boxes, making the C Boxes stateless. This offers an interesting benefit for those customers using SMB: in the event of C Box failures and/or reboots, SMB recovers very quickly (unlike most of its competitors that store the SMB state data in volatile DRAM). Permissions data is also kept in SCM, making recoveries much faster as well.

Features added in 2021 include support for NFS 4.1 (but not pNFS), an internal load balancer, asynchronous replication, user quotas, immutable snapshots, and a cloud-based dashboard that provides a single pane of glass for management of systems (which are all on premises). The dashboard includes an interesting feature that customers often comment on: it provides them the visibility to be able to easily identify data that can be deleted. In modern systems that are using a number of storage efficiency technologies (sometimes simultaneously), it can be difficult to identify data that can be safely deleted. VAST Data has resolved that issue.

VAST Data's go-to-market model brings interesting advantages to its customers. Although the vendor started off selling appliances, in 2021, it shifted to more of a software-defined model based around subscription pricing. This approach, dubbed "Gemini" by the vendor, leverages a single major distributor (Avnet) as the supplier of VAST Data's hardware (which includes a number of x86 server-based storage suppliers), while the vendor provides the software. Gemini is structured to preserve single SKU ordering, deliver appliances with preloaded VAST Data software, and offer customers a single point of support contact. But the approach has one other significant advantage for customers: it lets them buy hardware at a volume discount determined by all of VAST Data's customers' hardware purchases, a feature that effectively lets customers buy like a hyperscaler (with higher volume discounts than they could obtain on their own) but deploy like an enterprise (using appliances).

## **Challenges**

While VAST Data customers build very large production clusters (and in fact have consolidated multiple smaller clusters from other distributed scale-out file system platforms onto a single Universal Storage Platform), average response times tend to hover around a millisecond at scale. The vendor uses intelligent algorithms to maximize hit rates for data in the SCM, and latencies for that data are much lower than 1ms. Customers tend to consolidate a lot of workloads onto a VAST Data system, but they tend not to use it for their most latency-sensitive workloads.

VAST Data has optimized its EC (called Locally Decodable Erasure Codes) and data reduction (similarity-based data reduction) for systems that have over 1PB of raw capacity. While the economics of the Universal Storage Platform are very good at scale (1PB+), they are less so under that limit. Customers considering the system should probably have 1PB+ of data they will store on it (or will be getting there shortly after buying the system).

The Universal Storage Platform is not a software-defined platform to the same extent as some of its competitors. While the C Boxes are standard x86 servers running the VAST Data software in a container, the D Boxes are not (although they are an off-the-shelf product). Their unique design includes the PCIe switch and four capacity "nodes," and they can support up to 52 SSDs in 2U. This hardware is still delivered by Avnet (as are the C Boxes), but to date, they have only been available from one supplier (a potential cause for concern). VAST Data has been working to secure other sources for the D Box hardware, and in mid-2022, it will be introducing three additional sources. Even with the new sources, customers will continue to purchase VAST Data solutions from Avnet.

The system also requires a certain percentage of SCM capacity, which can be supplied by either Intel (Optane) or Kioxia (based on the FL6 design). And because of the particular configuration of this disaggregated architecture, customers cannot deploy a software-only version of the VAST Data OS in the public cloud. The vendor's software licenses are, however, transferable to new hardware if and when the customer decides to upgrade. And VAST Data enjoys another aspect of software defined: it supports asymmetric scaling that allows different device types and multigenerational nodes to be in the same cluster. Not all of its competitors support that flexibility.

### **Consider VAST Data When**

For enterprises with over 1PB of unstructured data they want to put on a single platform, VAST Data offers good performance and price/performance along with linear scalability and high infrastructure efficiency. It is much easier to deploy than many other distributed scale-out platforms because the networking infrastructure is simple and as an all-flash system dispenses with the headaches associated with HDDs at scale (uneven performance, reliability, rebuild times, HDD replacement effort, energy, and floorspace consumption). Among enterprises that have purchased VAST Data, the top workload use cases are analytics, AI, backup, cloud infrastructure (among its hyperscaler and cloud service provider customers), and high-performance computing. Note that the vendor also supports FIPS 140-2-certified encryption.

Because of VAST Data's DASE architecture, performance scales much more linearly than shared-nothing designs as there is very little east-west traffic. This combines with other features of the system – its all-flash media and its excellent capacity utilization across both data protection and data reduction – to create very compact systems to meet defined performance and capacity objectives. Don't evaluate the "cost" of a system up front by just looking at the blended cost per gigabyte media costs – when comparing vendors, look at the total cost of the system required to meet stated objectives. In many cases, the VAST Data system will be smaller and require less raw capacity, contributing to a lower cost.

### **Vendors to Watch**

Although it was not evaluated as part of this vendor assessment, VMware supports the creation of NFS and/or SMB datastores with clients built into ESXi. VMware does not market this capability separately, and it is not designed to deliver the kind of performance and scalability that products evaluated in this document are. While it can be used for enterprise file sharing, it is predominantly used to provide a

central repository under the control of VMware vCenter for ISO images, virtual machine templates, and other file data so that customers do not need to maintain separate file-based data outside the VMware environment. The vendor's approach to file-based support may change in the future, particularly as it evolves to support more composable infrastructure offerings through its Project Capitola.

There are a number of other cloud-based file storage solutions marketed to enterprises from vendors like CTERA, Nasuni, and Panzura. These offerings are targeted mostly at simplifying file sharing and file management in geographically distributed enterprises, and they typically replace a number of smaller file servers residing in on-premises locations, each of which must be managed separately. They all tend to require public cloud components, and in that way, they are different from the vendors assessed in this document whose products can be deployed on premises or in the public cloud as needed (cloud deployments are not required). One might also put Hammerspace among this group, although unlike the other vendors it provides an overlay that creates a unified namespace across both on-premises and off-premises deployment locations and does not require data to be moved into a proprietary scale-out file system (just the metadata must be moved).

Hyperscalers have also entered this space with their own native file-based offerings (which obviously require public cloud deployment). These offerings, such as Amazon EFS, Microsoft Azure Files, and Google Cloud Files, are typically targeted for traditional file sharing and are often considered as an alternative to on-premises, self-managed file servers distributed throughout a number of locations. These tend to be basic file-sharing services that lack enterprise-class capabilities, and it is why a number of the hyperscalers also offer other cloud-based file services based on enterprise-class storage OEM software like Amazon FSx (which can use NetApp ONTAP, OpenZFS, Windows File Server, or Lustre) or NetApp Azure Files (which is based on NetApp ONTAP). Providing enterprise-class storage services is a selling point for storage OEMs like Dell, NetApp, and Qumulo whose storage operating systems can be deployed in the public cloud – that is the point of offerings like Dell's OneFS for Google Cloud, NetApp's Cloud Volumes ONTAP, or Qumulo's Cloud Q. These "enterprise" offerings all claim to offer higher performance and scalability and enterprise storage management capabilities not found in "native" cloud file offerings from the hyperscalers. For certain file sharing workloads, however, all of these products may at times compete.

## APPENDIX

---

### Reading an IDC MarketScape Graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis, or strategies axis, indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScope represents the market share of each individual vendor within the specific market segment being assessed, not the overall storage-related revenue of the vendor.

Several suppliers offer different file system offerings, although they do not all necessarily compete in the distributed scale-out file system segment. In cases where the vendor offers two scale-out file system types, IDC has worked with the vendor to select the product that most closely fits within the inclusion criteria of this study.

## IDC MarketScope Methodology

IDC MarketScope criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailor the range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions on the IDC MarketScope, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in an effort to provide an accurate and consistent assessment of each vendor's characteristics, behavior, and capability.

## Market Definition

In July 2021, IDC introduced a new taxonomy for the file system market. There are four segments to the file system market: scale-up file storage, scale-up clusters, distributed scale-out file storage, and parallel scale-out file storage. The scale-up segment is small and shrinking in size, while all the growth is being driven by scale-out products. Briefly, scale-out file systems distribute data across nodes while presenting a single data access namespace. There are some differences, however, in how data is distributed between scale-up clusters and scale-out file storage. In scale-up clusters, data is rarely ever distributed across nodes, and the throughput to a given file is limited to the bandwidth of the single node from which it is served. In scale-out clusters, data in a single file can be distributed across nodes, a design which can improve access performance, data concurrency, and recovery time.

Scale-up clusters and distributed scale-out file storage routinely compete for the same business in enterprises, and this vendor assessment includes vendors from both segments. For more detail on how each of these segments is defined, see *Reclassifying File Storage – A New Approach for the Future of Digital Infrastructure* (IDC #US48051221, July 2021).

## Evolution in the Distributed Scale-Out File System Market

File system platforms have been widely used in the enterprise for file sharing. In the early 2000s, data under management grew, new types of file-sharing workloads emerged, and scale-out designs for file sharing began to become more popular. Distributed scale-out file systems became a mainstream alternative to the NetApp filers that dominated file sharing in the 1990s, and NetApp introduced a clustering capability to extend the scalability of its own offerings beginning with the release of "Clustered Data ONTAP" in the late 2000s.

Target workloads for these types of platforms included post-production and media streaming in the media and entertainment market, imaging and video, home directories, local and distributed file sharing, test and development, batch analytics, and backup/archive (although this latter workload was also a major target for many object-based storage vendors). Over most of the life of the scale-out file

system market, two platforms were clearly huge players in the market: Dell, which had obtained the Isilon (now PowerScale) product with the acquisition of EMC in 2016, and NetApp, which has been focused on enterprise file-based storage since the company's founding in 1992.

Over the past 10 years, what IDC refers to as "second generation," distributed scale-out file systems were introduced by a number of mostly start-up vendors (Cohesity, Huawei, Nutanix, Pure Storage, Qumulo, and VAST Data). These newer platforms were characterized by more software-defined designs, focused on providing easier management at scale, improved storage infrastructure efficiencies, and in general being more "cloud-friendly." Some were specifically designed around newer storage technologies like NVMe, storage-class memory, and NVMe over Fabrics (NVMe-oF). Both Dell and NetApp have responded, and the distributed scale-out file system market is very different in 2022 from what it was in 2012.

Today, most vendors claim performance and scalability as differentiators, but in selecting a platform, enterprises should focus more on the different facets of performance and scalability that are important to their workloads as there are significant differences in these capabilities across the vendors reviewed in this study. Ease of use is another major differentiator between vendors. Selecting the right scale-out file-based storage platform demands that potential buyers look beyond high-level marketing messages proffered by vendors to understand which products best fit their unique storage I/O requirements.

### ***Key Differences Among Vendors in Product Design Strategies***

In today's digitally transforming world, enterprises are capturing, storing, protecting, and analyzing more data than ever before to drive better business insights in much more data-centric business models. To accommodate newer big data analytics workloads and increased scale, roughly 70% of enterprises going through digital transformation also plan to modernize their server, storage, and/or data protection infrastructure by 2023. In doing so, they are looking for more deployment and purchasing flexibility, simplified management at multi-petabyte levels of scale, increased performance and availability, better affinity with a hybrid multicloud world, and improved infrastructure efficiencies that allow them to pursue denser storage workload consolidation (to narrow not only the number of storage platforms that must be supported but also the number of vendors).

There are several areas where the designs and product strategy focus of certain vendors diverge:

- **Software-defined storage.** "Software defined" is all about improving flexibility, whether that is the flexibility to deploy on different types of server hardware from different vendors, the ability to deploy the file system stack in the public cloud, or the ability to easily accommodate new storage devices and technologies over time. Software defined also tends to offer better technology refresh models. While most enterprises want to buy appliances that offer a single point of support contact and single SKU purchasing, they like the ability to select the hardware of their choice and have a vendor deliver it that way (usually through a channel partner that creates the combined hardware/software platform). Some vendors only deliver appliances on a single type of server hardware (i.e., their own), while others offer a variety of hardware options with their appliances.
- **Access methods.** Most of the assessed vendors are committed to using traditional access methods like NFS and SMB that do not require the installation of custom software on the client side. While there are certain performance scalability limitations to the use of native protocols, they do tend to meet most of the needs of enterprise workloads, and they are easy to deploy and manage. A number of the assessed vendors support options (like "nconnect" for NFS) to improve the scalability of native access methods. Parallel scale-out file systems (which are not

being assessed in this document but will be assessed in a future IDC MarketScape document) use proprietary intelligent clients that support a parallelism that allows their throughput to a single large file to go significantly beyond where NFS and/or SMB can go, but a few legacy enterprise workloads can benefit from that. Many enterprises are deploying AI-driven big data analytics workloads as part of digital transformation, and certain stages of the AI data pipeline actually can benefit significantly from this increased throughput though. Several of the assessed vendors not only focus on native access methods (e.g., NFS and SMB) but also offer an intelligent client option that enterprises can deploy and use if/when needed.

In general, the more access methods a scale-out file-based storage platform supports, the more options there are for denser workload consolidation. While NFS and SMB are the most popular file-based access methods, a number of vendors support other options as well like FTP, HTTP, and HDFS (although many HDFS workloads that are being modernized are moving to object-based storage). Many new applications being developed and deployed during digital transformation use Amazon's Simple Storage Service (S3) interface. Even though that is an object-based access method, an increasing number of vendors allow their file-based storage to be accessed over S3.

- **Cloud native.** Many new workloads are deployed in the public cloud, and enterprises are at the same time evaluating the disposition of their legacy on-premises workloads (rehost into a virtual machine [VM], refactor for cloud deployment, re-architect for optimized cloud deployment, replace [usually with a cloud-based version], or retire). More software-defined designs make this easier since a software-only product can also be deployed on web-scale infrastructure in a public cloud environment. Other features impacting "cloud-friendliness" include microservices design, container-based deployment, support for the Container Storage Interface (CSI, an interface that allows the storage system to provide persistent storage to applications running in containers), APIs that support Kubernetes-based automation, and unified management consoles that provide comprehensive visibility into workloads that span on-premises and off-premises deployment models (e.g., an instance of a distributed scale-out file system running on hyperconverged infrastructure [HCI] on premises and also on web-scale infrastructure in a public cloud, with the two instances collaborating on a workflow).

Subscription-based licensing may also be viewed as a "cloudlike" capability, and most of the vendors offer this as an option (or as the only way to purchase their software). Several of the vendors (Cohesity, Nutanix, Qumulo, and VAST Data) started out as appliance vendors but have moved to a new software-only business model, making appliances available through channel partners. As long as customers still enjoy the ability to buy appliances that ease purchasing and deployment and have a single point of support contact for the file-based storage solution, they tend to be indifferent to this change, but it can have a major impact on vendor operations, margins, and revenue per employee.

- **Unified storage or multiprotocol access.** There are several different product strategies here. Two of the vendors (Huawei and Nutanix) can support block-, file-, and object-based storage on a single system, all managed from a unified interface. The software-defined nature of both of these platforms provides the flexibility to configure different storage pools within the system for different I/O profiles and access methods. One vendor supports both file- and object-based access methods (but not block) to data, although the data organization method is a key value store (Pure Storage). (Note that Huawei also supports this approach for file and object data, but it uses a separate volume-based data organization method for block.) If a data object can be natively written using either a file-based interface or an object-based interface to the key value store but to have it natively accessible by multiple interfaces requires multiple copies of the data, IDC refers to that as "unified storage."

Other vendors use a file-based data organization method but support multiprotocol access to the same data through a variety of interfaces like NFS, SMB, HDFS, FTP, HTTP, NDMP, and S3. (IDC refers to this as "multiprotocol access.") With multiprotocol access, storage capacity is used more efficiently (both NFS and S3 access the same underlying data object), but the issue of semantic loss may arise. Semantic loss occurs when the interface through which data is accessed (e.g., NFS) does not support all the primitives of the interface used to initially write the data (e.g., S3) potentially limiting an application trying to access non-native data. This can be an issue for some applications but not for others.

- The ability to support multiple access methods to the same data on a single platform can make working with multistage data pipelines much easier. Data does not need to be copied over a network to another system, which must be managed separately and may require a different administrative skill set. Fewer systems can be purchased, and potentially fewer vendors can be managed. Software-defined flexibility can allow storage to be configured in a single system to meet a variety of different I/O profiles that may be required in different stages of a data pipeline. Sharing data can allow workload consolidation onto fewer platforms, but there are clearly caveats in doing so (risks to meeting performance and/or availability SLAs, security in sharing data across workloads, etc.).

- **On-disk data protection.** While file systems historically used replicas to protect data on disk, vendors have introduced interesting innovations that can provide better capacity utilization and higher durability and/or enable higher performance access to files. EC, which had historically been used in object-based storage platforms, is now available on a number of the vendor offerings evaluated in this assessment. EC distributes data more widely across devices and/or nodes, splitting it into data and parity bits. This approach makes better use of available storage capacity (than making full file copies) while offering the same or better data durability. Like replicas, EC can span geographical sites to provide site-level resiliency.

One vendor (VAST Data) has implemented an EC approach that offers significantly lower capacity utilization to meet data resiliency requirements for at-scale configurations (larger than 1PB) by distributing data more widely than any other vendor. Other vendors can use replicas for files below a certain size, and then transparently switch to EC as files get larger to save space. (Use of replicas for data protection may produce lower access latencies for small files, while more widely distributed EC can provide higher throughput for larger files.) And other vendors use a more RAID-like approach that operates at the block level, which allows it to recover just the missing data rather than full files for faster rebuilds. Each approach has its pros and cons, depending on customer preference for access latency, fast rebuild times, lower capacity utilization, and data durability.

- **Storage architecture.** All of the assessed systems but one (Pure Storage FlashBlade) use cache-based architectures. As write ingest scales, cache-based architectures can eventually hit a "write cliff" when the write cache is not being destaged to persistent storage as fast as it is filling up. When this occurs, write performance drops noticeably. The write cliff can be pushed out by adding nodes and more widely distributing the load (which all vendors do), creating larger high-speed caching tiers (which VAST Data has done by using a very large persistent storage-class memory-based layer as a write cache), and using other innovative software-based techniques to extend write performance. Cacheless architectures write directly to persistent storage, which is of course much larger than any caching tier (persistent or volatile) but often offers higher write latencies because persistent storage usually is slower than the memory media generally used in caching tiers. Pure Storage, which uses only solid state media in its FlashBlade, will explain why it thinks its approach is a better fit for file- and object-based workloads in the enterprise, and customers can decide for themselves which is best for their environment.

- **Data management strategies.** As rapid data growth continues, enterprises want features that allow them to implement more efficient data management strategies. Although data tiering (both within a system and to external targets) has been available for a long time, providing the visibility that enables effective data classification is really the right starting point for intelligent data management. While file usage metrics like frequency of access have long been used to determine data placement, new AI and machine learning (ML) technologies allow data placement to be better optimized in real time for performance as well as across tiers to reduce infrastructure and management costs. They can also identify data that can be safely deleted, ensuring that enterprises are only keeping data that has to be retained. While some vendors realized the importance of intelligent data placement early on, at this point, pretty much all of them are doing at least something in this area beyond just tracking frequency of access.
- **Deploying the scale-out file system.** All of the vendors require data to be migrated from third-party file systems into their own scale-out file system design. This is considered standard operating procedure to be able to take advantage of all the features of distributed scale-out file systems but is in contrast to certain file-based storage players (not evaluated in this vendor assessment) whose software layers on top of existing file systems to provide a unified namespace (requiring no data movement). One vendor in this assessment uses a very different strategy, enabling backup data to be converted into one or more scalable, shared access scale-out file systems. The strategy of this vendor (Cohesity) effectively combines data protection and enterprise file sharing into a single, centrally managed platform. Some enterprises find this very convenient, while others prefer to keep the two practices (data protection and file management) separate.

What enterprises can expect from half of the players in this market is that the scale-out file system software runs on commodity x86 hardware, and vendors generally support multiple hardware options (e.g., Dell PowerEdge, HPE ProLiant, Lenovo, and/or Supermicro). Four of the vendors still prefer that customers buy appliances running their hardware of choice (Dell PowerScale, Huawei OceanStor Pacific, NetApp ONTAP, and Pure Storage FlashBlade), claiming that there are performance, availability, and/or other advantages that accrue. IDC expects that in the future most successful vendors in this space will provide a software-only version that can run on commodity x86 hardware, if for no other reason than to offer the opportunity to deploy their file system in public clouds (several vendors in this study in fact already do this). But it is true that many enterprises see advantages to using purpose built versus commodity hardware – Dell, Huawei, NetApp, and Pure Storage do this and are huge market players in scale-out file systems, although revenue growth rates are higher among the start-up players.

Other baseline expectations should include entry-level configurations that are highly available to require at least three nodes (although non-highly available configurations for edge deployments may be supported in a single VM), an ability to scale performance and capacity by adding nodes up to 100+ (in some cases quite a bit more), and an ability to mix node types (e.g., performance-intensive nodes that might be all-NVMe, hybrid nodes that can support a mix of SSD and HDD, and archive or capacity-intensive nodes that may be all HDDs). This ability to support mixed nodes enables most of these systems to offer a technology refresh model that is nondisruptive (just add the new node types as they become available) and can preserve existing investment (customers are not required to get rid of older nodes if they don't want to even as they add newer nodes).

## Strategies and Capabilities Criteria

This IDC MarketScape evaluates vendors by analyzing both their current capabilities and their strategies going forward. It uses a scoring and assessment model based on both qualitative and

quantitative criteria, which results in a single graphical illustration (refer back to Figure 1) that shows vendors' positions in the market relative to each other. These positions are driven by architectural and other considerations that are discussed in this study.

Tables 1 and 2 match the two axes of the IDC MarketScape Figure (refer back to Figure 1), with Table 1 covering strategies and Table 2 covering capabilities. Tables 1 and 2 provide a definition of the applicable criteria with the weighting applied to each. The weighting for each set of criteria is in a percentage, and all of the weightings for each set (strategies and capabilities) sum to 100%. A higher weighting indicates that a specific criterion is more important than a criterion with a lower weighting.

**TABLE 1**

**Key Strategy Measures for Success: Worldwide Distributed Scale-Out File System**

Strategies Criteria	Definition	Weight (%)
Cloud-native workload support strategy	Microservices-based storage OS, storage OS runs in a container, support for container-based persistent storage, and so forth	10.00
Consumption models	Appliance, software only, subscription, managed service, converged infrastructure, reference architectures, and so forth	12.00
Cost	Infrastructure efficiency, TCO, cost to acquire/deploy/expand/upgrade/refresh, transferable licensing, all-inclusive software, and so forth	15.00
Customer experience (CX)	Simplifying purchases, guarantees, trade-in credits, NPS, CX programs, and so forth.	7.50
Marketplace creation	Strategy for creating an Apple iPhone like ISV marketplace with APIs for custom development, VMware support (Tanzu), Red Hat OpenShift, Kubernetes, accelerated compute, and partnering strategy	2.50
Functionality or offering	Block, file, and/or object	3.00
Market traction	Overall revenue, recent growth rates, installed base and customer count, revenue diversity, OEM agreements, and so forth	15.00
Road map	Areas of focus going forward, access methods, cloud support, software defined, data classification/mobility, and so forth	10.00
Tech support coverage	Geographic coverage and coverage options	10.00
Top differentiators	Judged sustainable differentiators around performance, scalability, availability, flexibility, ease of use, cost, and so forth	15.00
Total		100.00

Source: IDC, 2022

**TABLE 2**

**Key Capabilities Measures for Success: Worldwide Distributed Scale-Out File System**

Capabilities Criteria	Definition	Weight (%)
Availability/recovery	Inline data protection methods, one-click failover/failback, one-click restore, snap to object, ransomware (air gap), replication, NDU, SMB failover, data rebalancing, retiring old nodes, and so forth	10.00
Security/ransomware	Encryption, FIPS compliance, RBAC, air-gap protection (immutability), tape support, other ransomware detection/prevention/recovery, and so forth	9.00
Cloud support	Public cloud support/relationships, first-party offerings, cloud-based deployment, cloud SLAs, unified management/billing, and so forth	5.00
Data reduction	Compression, deduplication, hardware assist, additional post-process data reduction, space-efficient snapshots, delta differentials, write coalescing, and so forth	8.00
Ease of use	Use of defaults versus having to manually configure everything, making things easier to manage as they scale, one-click operations, and so forth	12.50
Flexibility/agility	Mixed node and media support; deploy on bare metal, VMs, containers, and in the cloud; how software defined is it; and life cycle/tech refresh	10.00
Functionality or offering	Scalability and data services	7.50
Multiprotocol support	Which access methods, unified or separate, semantic loss, NFS, SMB, HDFS, FTP, HTTP, S3, SWIFT, CSI, NDMP, S3 support strategy, intelligent client option, and so forth	9.00
Performance	Serial ingest, metadata-intensive workloads, random read/write against small files, large file reads, latencies, throughput, GPUDirect Support, NVMe, and so forth	10.00
Scalability	Entry and maximum configurations, back-end network congestion, global lock management, edge deployment option, max file/file system size, small file optimizations, and largest production cluster size	10.00
Snapshots	Read only (immutable), read/write, space efficient, what kind (ROW or COW), how many snapshots, and simple snap to object	7.00
Unified namespace	Unifies file and object in the same or across different platforms, between on-prem and public cloud-based deployments, don't really support object, and so forth	2.00
<b>Total</b>		<b>100.00</b>

Source: IDC, 2022

## LEARN MORE

---

### Related Research

- *Qumulo Exhibiting Strong Momentum as It Serves the Evolving Unstructured Data Storage Needs of Enterprise Customers* (IDC #US48896622, March 2022)
- *VAST Data: A Technical Deep-Dive Look at a Compelling New Scale-Out Storage Architecture* (IDC #US48805222, February 2022)
- *Worldwide File- and Object-Based Storage Forecast, 2021-2025: New Enterprise Workloads Driving Strong Growth* (IDC #US48403021, December 2021)
- *Hyperconverged Infrastructure Adoption Trends – 3Q21: Building Block for Hybrid Cloud Infrastructure* (IDC #US48308121, October 2021)
- *Enterprise Workloads Resulting in Broader Adoption of Scale-Out File Storage Architectures* (IDC #US48305121, October 2021)
- *How to Compare Distributed Scale-Out File Storage Platforms for Use with Enterprise Workloads* (IDC #US48191621, September 2021)
- *Reclassifying File Storage – A New Approach for the Future of Digital Infrastructure* (IDC #US48051221, July 2021)

### Synopsis

This IDC study represents a vendor assessment model called the IDC MarketScope. This study is a quantitative and qualitative assessment of the characteristics that assess a vendor's current and future success in the relevant market or market segment and provide a measure of the vendor's ability to become a leader or maintain leadership.

The distributed scale-out file system market segment, which is part of the file-based storage market, is an example of a large, maturing market that is still exhibiting low double-digit growth. This document assesses the capabilities and strategies of key vendors of scale-out file-based platforms. While seven of the assessed vendors have distributed scale-out file system designs, one of the vendors (NetApp) actually uses a scale-up cluster design but still meets the inclusion criteria of this vendor assessment study.

"While all evaluated vendors tout the performance, scalability, and ease of use of their file-based storage offerings, a closer look reveals important distinctions in how vendors define these metrics and build their products to achieve them," said Eric Burgener, research vice president, Infrastructure Systems, Platforms and Technologies Group, IDC. "To select the right product, enterprises need to understand the architectural differences between the different vendor approaches, understand the implications of those choices for their workloads, and then choose the product which best fits their requirements."

## 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 and Trademark Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit [www.idc.com](http://www.idc.com) to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit [www.idc.com/offices](http://www.idc.com/offices). Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or [sales@idc.com](mailto:sales@idc.com) for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights. IDC and IDC MarketScape are trademarks of International Data Group, Inc.

Copyright 2022 IDC. Reproduction is forbidden unless authorized. All rights reserved.

