

Magic Quadrant for Distributed File Systems and Object Storage

Published 1 October 2021 - ID G00738148 - 39 min read

By Julia Palmer, Jerry Rozeman, [and 2 more](#)

IT leaders are looking to deliver cost-effective, scalable, flexible storage platforms for growing amounts of unstructured data. This research helps I&O leaders to assess the key attributes, vision and execution of distributed file systems and object storage vendors.

Strategic Planning Assumptions

By 2026, large enterprises will triple their unstructured data capacity stored as file or object storage on-premises, at the edge or in the public cloud, compared to 2021.

By 2025, 40% of infrastructure and operations (I&O) leaders will implement at least one hybrid cloud storage capability, up from 15% in 2021.

By 2025, 60% of the global unstructured data storage capacity will be deployed as software-defined storage (SDS), up from less than 25% in 2021.

Market Definition/Description

The unstructured data storage market is rapidly evolving. New capabilities are needed to address the growing challenges of exponential data growth, rapid digitalization and the globalization of business. Storage infrastructure products, based on distributed file systems and object storage, are increasing faster than ever in both volume and capacity of deployments as the platform of choice to address the growth of unstructured data in global data centers.

Gartner end users report unstructured data growing more than 30% year over year. Thus, I&O leaders are looking for extensible on-premises storage products that can address an increasing number of digital business use cases with lower acquisition, operational and management costs. I&O leaders are demanding scalability, flexibility, life cycle management, ease of management and analytics insights into data. In recent years, Gartner clients are also asking for cyber-resilient unstructured data solutions to prevent, detect and recover from ransomware attacks.

The steep growth of unstructured data for emerging and established workloads now requires new types of products and cost efficiencies. Most products in this market are driven by SDS, capable of delivering tens of petabytes of storage. SDS can also potentially leverage hybrid cloud workflows with public cloud IaaS to lower total cost of ownership (TCO) and improve data mobility. New and established storage vendors continue to develop scalable storage clustered file systems and object storage products to address cost, agility and scalability limitations in traditional, scale-up storage environments.

Gartner defines distributed file systems and object storage as software and hardware platforms that are based on distributed designs and that support object and/or scale-out file system technology to address requirements for unstructured data growth. This market is based on distributed computing architectures in which there is no single point of failure or contention across the system. More specifically, products must have a fully distributed architecture where data and metadata are distributed, replicated or erasure-coded over the network across multiple nodes in the cluster. It's essential, when managing multi-petabyte-scale systems, to have support for horizontal scaling of capacity and throughput in a cluster through an independent node addition to a global namespace/file system.

Two underlying technologies that form the market for unstructured data storage are defined as:

- **Distributed file system storage**, which uses a single parallel file system to cluster multiple storage nodes together, presenting a single namespace and a storage pool to provide high-bandwidth data access for multiple hosts in parallel. Data and metadata are distributed over multiple nodes in the cluster to deliver data availability and resilience in a self-healing manner, and to scale capacity and throughput linearly.
- **Object storage**, which refers to systems and software that house data in structures called “objects” and serve clients data via RESTful HTTP APIs, such as Amazon Simple Storage Service (S3), which has become the de facto standard for accessing object storage.

Magic Quadrant

Figure 1: Magic Quadrant for Distributed File Systems and Object Storage





Source: Gartner (October 2021)

Vendor Strengths and Cautions

Cloudian

Cloudian is a Challenger in this Magic Quadrant. Cloudian offers a scale-out object platform that is designed for high-throughput object storage workloads. Cloudian HyperStore runs on-premises and in public clouds. It offers an add-on scale-out file gateway to handle file workloads. Its operations are mostly focused in North America and EMEA, and its clients tend to be enterprise customers across the public sector, service providers and large enterprises. Over the past 12 months, Cloudian added several improvements such as support for Kubernetes, VMware Tanzu and all-flash object storage, and improved its file gateway solution, HyperFile, to a scale-out architecture.

Strengths

- Cloudian's strategy to focus on backup opportunities for providing S3-compatible storage with Object Lock technology for ransomware protection remains an attractive option for many customers.
- Customers express a high degree of satisfaction with the HyperStore product for its simplicity of deployment, ease of management, S3 compatibility and quality of support.
- Cloudian's flexible software-based and hardware-agnostic deployment options, branded appliances, and cloud-based deployment are attractive options for many customers.

Cautions

- Cloudian's global installed base is relatively small compared to other market leaders, and it has a minimal installed base in the Asia/Pacific region.
- Cloudian's file service implementation, HyperFile, is not natively integrated into the object storage system as it is an add-on gateway solution.
- Cloudian has a limited presence in use cases apart from backup and archiving and, as such, IT leaders should make a broader evaluation for other use cases.

DDN

DDN is a Niche Player in this Magic Quadrant. DDN EXAScaler is a distributed file system that runs on-premises and in the cloud. It is deployed primarily for large-scale, high-throughput file workloads and, in a few cases, as an S3 target for archive use cases. EXAScaler is powered by the Lustre open-source parallel file system. DDN's operations are geographically diversified and its clients tend to be large enterprises. Over the past 12 months, DDN delivered enhancements to EXAScaler's protocol services and optimized cloud-native deployment on the three major public clouds. In addition, DDN gained certification for deployments in the public cloud with SAS Grid Manager, a grid computing manager. DDN EXAScaler is best suited for high-performance computing (HPC) and analytics use cases.

Strengths

- DDN offers a proven storage solution for HPC and artificial intelligence (AI) workloads, including enhancements for deployment and integration in GPU environments to support NVIDIA DGX SuperPOD.
- DDN has been shipping enterprise-grade storage products for more than a decade, and has established and matured operations.
- DDN's parallel file system is now available in all three major public cloud IaaS providers, giving customers the choice to extend their applications to any of those public clouds.

Cautions

- DDN's offerings have limited focus in addressing requirements for general-purpose file storage.
- DDN does not have developer-friendly object storage for object-native workloads that is simple to download, deploy and manage.
- DDN's file offering lacks robust AIOps-based monitoring and management, a requirement for large-scale and performance-sensitive workloads, making it complex to operate and troubleshoot.

Dell Technologies

Dell Technologies is a Leader in this Magic Quadrant. Dell EMC PowerScale is a distributed file system that runs on-premises and is available as a native service in Google Cloud (Dell EMC PowerScale for Google Cloud). Dell EMC ECS is an object storage platform that runs on-premises and in the cloud. Its operations are geographically diversified and its clients range from small to very large enterprises across all verticals. Over the last 12 months, Dell EMC PowerScale added support for NVIDIA GPUDirect, parallel upgrades, in-line compression and data deduplication. Dell EMC ECS added support for larger SSDs and HDDs, S3 access, an external key manager, and object tagging. Both products cover the majority of the unstructured data use cases and offer centralized management with Dell EMC CloudIQ and Dell EMC DataIQ.

Strengths

- Dell Technologies is a trusted brand name for data storage worldwide and it has the largest installed base of unstructured data storage with PowerScale and ECS.
- Dell Technologies' file and object storage products are built on the same proven Dell server hardware, giving clients an additional level of assurance on product reliability and parts availability around the globe.
- Dell EMC PowerScale offers support for a wide variety of use cases and deployment options, such as GCP native file service and an option to consume as storage-as-a-service through APEX.

Cautions

- Dell Technologies lacks an SDS offering that can run on non-Dell bare-metal x86 servers and lacks market traction in public cloud deployments, resulting in limited applicability as a hybrid cloud storage solution.
- Relative to other products in the market, prospective customers have called out the pricing as high for both PowerScale and ECS.
- Dell Technologies lacks a common platform strategy for addressing both file and object native workloads. PowerScale and ECS remain separate products, with little integration beyond tiering from PowerScale to ECS.

Hitachi Vantara

Hitachi Vantara is a Challenger in this Magic Quadrant. Hitachi Vantara's Hitachi Content Platform (HCP) is an object storage product offered as software or an appliance that runs on-premises, and in public and hybrid cloud environments. HCP can scale performance and capacity independently to support a wide range of workloads. Hitachi Vantara operations are in the Americas, EMEA and the Asia/Pacific region. Clients tend to be midsize to large enterprises. Over the past 12 months, Hitachi Vantara has added the following capabilities: OEM integration with WekaFS; EverFlex for XaaS, and HCP Anywhere file sync and share capability integrated with HCP object storage. HCP is best suited for analytics, cloud storage, backup and archive, and hybrid cloud.

Strengths

- Hitachi Vantara HCP is trusted for its reliability, and has an extensive and loyal global customer base that favors its high availability and resilient systems.
- Hitachi Vantara's HCP portfolio and analytics for data management have an extensive ecosystem that spans a wide range of verticals and solution domains.
- Hitachi Vantara has developed HCP integration with public clouds (Amazon Web Services [AWS], Microsoft Azure and Google Cloud Platform [GCP]) for data tiering and mobility.

Cautions

- According to Gartner clients, Hitachi Vantara is one of the more difficult vendors among industry leaders to work with and can be inflexible with contractual matters.
- Hitachi Vantara lags behind industry leaders in the use of real-time data collation in support of back-end customer care capabilities for timely service issue resolution.
- Hitachi relies on Weka's WekaFS software and integrated OEM solution partnership, as one of its file systems, to address performance-intensive, file-heavy workloads, so clients should carefully consider Weka as part of this solution.

Huawei

Huawei is a Challenger in this Magic Quadrant. Huawei offers OceanStor Pacific (previously OceanStor 100D) as a single distributed file system, block and object storage product. While OceanStor 9000 is still supported for file workloads, OceanStor Pacific is now Huawei's flagship product for all unstructured data needs. It has large operations in the Asia/Pacific region, and its clients tend to be telecom, government and financial institutions. Over the past year, OceanStor Pacific added new high-density capacity and performance nodes. It now has asynchronous replication, a distributed parallel client, multitenancy and encryption, as well as compression and Active Directory and LDAP authentication. OceanStor Pacific is best suited for large private clouds, content distribution, cloud-native applications and archiving.

Strengths

- Huawei is an established and trusted storage provider of file and object solutions with broad adoption in the financial, media and government verticals.
- Huawei's OceanStor Pacific is well-designed to accommodate a wide range of unstructured data workloads as a single product for scalable and performant file and object storage services.
- OceanStor Pacific has some of the best data efficiency capabilities in the market.

Cautions

- U.S. sanctions and geopolitical impact will continue to challenge Huawei's adoption and expansion strategy for enterprises outside the Asia/Pacific region.
- OceanStor Pacific file services are still nascent in terms of feature set and customer adoption.
- Huawei lacks thorough installation and configuration documentation for complex implementations, which may result in longer deployments and the need for vendor support and services.

IBM

IBM is a Leader in this Magic Quadrant. IBM Spectrum Scale is a parallel file system product that runs on-premises and in select public clouds. IBM's Cloud Object Storage (COS) is an object storage offering that runs on-premises and in the IBM Cloud. Its operations are geographically diversified and its clients range from midsize to very large enterprises. Over the past 12 months, Spectrum Scale enhanced support for hybrid cloud storage, containerized applications, file auditing and QoS. IBM COS added File Access, a virtual edge caching solution that provides a standard file protocol (SMB and NFS) interface. Spectrum Scale is best suited for commercial HPC and analytics use cases. COS is best suited for archive and private cloud storage.

Strengths

- IBM's acquisition of Red Hat expands its reach into enterprises with cloud-native workloads, including enhanced support for containerized applications and OpenShift in IBM Spectrum Scale.
- IBM COS is the only object storage offering on the market that is also the underlying storage of a major public cloud. This provides the assurance of running a large-scale environment for object native workloads.
- IBM Spectrum Discover analyzes data stored in IBM Spectrum Scale and COS, providing data visibility, classification and labeling with custom metadata to enable AI-powered applications.

Cautions

- IT leaders rarely consider Spectrum Scale for general-purpose file system workloads because it was designed for high-performance workloads and depends on open-source NFS and SMB implementations.
- IBM's offerings lag behind leading competitive offerings in AIOps-powered troubleshooting and monitoring capabilities.
- End users have expressed that IBM's COS user interface is complex and not intuitive or user-friendly.

Inspur

Inspur is a Niche Player in this Magic Quadrant. The Inspur AS13000G5 series provides a unified software solution for both file and object storage. Inspur offers three AS13000G5 models for petabyte-scale applications for high performance, high-definition video, high reliability and cloud-based deployments. Its operations are mostly focused in the Asia/Pacific region and EMEA, and its clients tend to be telecom, government, internet, mobile and education institutions. Over the past 12 months, Inspur has added two new arrays, tagging, native HDFS, two-way authentication of HTTPS, S3 select, capacity monitoring and predictive analytics. Inspur storage is best suited for backup and archiving, commercial HPC, hybrid cloud, and analytics.

Strengths

- Inspur has experienced a rapid growth in customers and managed petabytes in the China and Japan regions, due to its product features and strong position in the server market.
- Inspur receives high marks from its customers for aggressive pricing, high performance and large cluster capacity – three factors that drive its competitive advantage in the Chinese market for HPC and backup/archive.
- Inspur has made significant investments in its AIOps-based InView product to predict disk failures, capacity and performance issues before they occur.

Cautions

- Geopolitical implications to U.S. sanctions outside of the Asia/Pacific region may hinder Inspur's market expansion initiatives.
- Inspur lags behind major industry vendors with a proven centrally managed, SaaS-based storage-as-a-service offering for data services across major cloud platforms outside of the Asia/Pacific region.
- Inspur's AS13000G5 capabilities and features are most focused on the performance-oriented use cases, which will limit its applicability as a single unstructured data platform.

NetApp

NetApp is a Visionary in this Magic Quadrant. NetApp StorageGRID is an object storage solution available as software and hardware appliances that can run on-premises and in the public cloud. NetApp supports tiering of data from on-premises StorageGRID to public cloud services, including AWS and Azure. NetApp operations are global, and its clients tend to be large enterprises, media and entertainment (M&E), government, and service providers. Over the past 12 months, NetApp added enhancements for the security, performance and monitoring of StorageGRID. It also added fully managed object storage services in partnership with Equinix. NetApp is best suited for cloud storage, archiving, backup and hybrid cloud.

Strengths

- NetApp offers a broad portfolio of appliances ranging from cost-optimized dense platforms to performance-focused appliances for transactional object workloads.
- NetApp's StorageGRID is frequently used to deploy multiple workloads, including both traditional backup/archive and modern object native primary storage workloads.
- NetApp's fabric pool capability works in concert with Active IQ to provide easy tiering of files from NetApp FAS systems into StorageGRID.

Cautions

- NetApp lacks a common unstructured data platform for addressing both distributed file and object native workloads.
- Sales and customer support for NetApp StorageGRID gets secondary focus relative to NetApp's larger and more successful products, such as SSA and public cloud.
- NetApp users have expressed that StorageGRID's GUI and UX can benefit from usability enhancements.

Nutanix

Nutanix is a Visionary in this Magic Quadrant. Nutanix Files and Nutanix Objects are integrated with the Nutanix hyperconverged platform to provide a single architecture and management. It is designed to be deployed across nodes, leveraging an existing hyperconverged infrastructure (HCI) for smaller deployments, or on a dedicated-storage-only cluster for larger ones. Nutanix operations are global and rely on server OEM and channel partners for solution delivery. Nutanix has customers across all verticals. Over the last year, Nutanix Files added ransomware detection, improved disaster recovery, enhanced file analytics and synchronization capabilities. Nutanix Objects has improved performance and replication, and added tiering to S3 storage. Nutanix Objects also added smart lock, legal hold, and versioning and multitenancy support. Nutanix is best suited for hybrid cloud and cloud-native application storage.

Strengths

- Nutanix Files and Nutanix Objects offer existing HCI users an easy way to enable enterprise-grade software-defined distributed file or object services at any scale.
- Nutanix's rich data analytics tools are providing differentiating capabilities such as monitoring and reporting of abnormal user behavior, performance anomalies and audit trails, as well as built-in ransomware protection with the ability to detect and block attacks.
- Gartner clients highlight product ease of use and simplicity of management compared to the competition.

Cautions

- Nutanix Files and Nutanix Objects are not often considered by Gartner clients as stand-alone storage products, which has historically limited their applicability to only target existing Nutanix HCI customers.
- Gartner clients highlight that Nutanix Files might not be cost-effective for long-term unstructured data archiving.
- Nutanix Files and Nutanix Objects do not support deduplication and their public cloud presence is limited to Nutanix clusters on AWS bare metal.

Pure Storage

Pure Storage is a Leader in this Magic Quadrant. Pure Storage FlashBlade is a purpose-built unified file and object platform. Pure Storage offers a scale-out distributed file system that is designed to handle tens of billions of files and objects for massive throughput and parallelism by simply adding blades to scale capacity and performance. Pure Storage FlashBlade's primary market adoption has been in North America. Over the past 12 months, the vendor has added a number of enhancements, including native support for fast scale-out SMB 2.1, support for NFSv4.1 Kerberos and ACLs, enhanced NFS and object features, and Active Directory security with Kerberos authentication. Pure Storage is best suited for commercial HPC, analytics and backup where recovery time objective (RTO) performance is critical.

Strengths

- The vendor's strength in its FlashArray business and common management system provides its FlashBlade business with an advantage for customers that want to consolidate vendors across block, file and object workloads.
- Pure Storage's investments in its Pure1 AIOps and analytics capabilities provide its product with a compelling advantage in terms of a more resilient platform and proactive customer support.

- Pure Storage is well-regarded by its clients as easy to do business with, and its product is seen as simple to install and manage. Its Evergreen support, as part of a life cycle management program, is also highly valued.

Cautions

- Pure Storage FlashBlade capabilities and features are most suitably aligned to performance-oriented workloads, potentially excluding broad adoption as a single platform for all unstructured data workloads.
- Pure Storage lacks advanced enterprise features such as SMB 3.0 and seamless failover, deduplication, tiering, and the ability to deploy as a software-only version on public cloud infrastructure.
- Pure Storage's sole focus on an all-flash appliance makes the economics of its solution unsuitable for archive and long-term repository use cases that are served by hybrid arrays.

Quantum

Quantum is a Visionary in this Magic Quadrant. Quantum's ActiveScale is an object storage appliance that was acquired from Western Digital in March 2020. Quantum operates globally with a focus on North America and EMEA. Over the past 12 months, Quantum integrated the ActiveScale organization, and created a roadmap and vision for integrating ActiveScale into the portfolio by developing an integrated solution together with its tape and StorNext file system product line. Quantum also introduced subscription-based pricing models and acquired CatDV for metadata tagging of data in the ActiveScale product. Quantum ActiveScale is best suited for large unstructured data workloads such as archiving of scientific and medical research data and repositories of rich media content.

Strengths

- Quantum is a respected and attractive vendor for organizations looking for end-to-end solutions that store high volumes of rich media such as video data.
- Quantum's vision differentiates by integrating several technologies such as its StorNext file system, ActiveScale object storage and tape solutions into a single platform to deliver an eternal archive, and will attract I&O leaders looking for cost-optimized, long-term data management and archiving.
- The vendor's CatDV acquisition will allow Quantum to add rich data services, enabling organizations to index, catalog and search for content in ActiveScale technology.

Cautions

- ActiveScale is a relatively new product in Quantum's portfolio and has experienced slow market adoption since the Western Digital acquisition. IT leaders should monitor growth and customer adoption to be sure of product continuity.
- Quantum's ActiveScale strategy is mainly around backup and large unstructured repositories, with a focus on the unlimited and eternal archive, and it is rarely considered for other use cases by Gartner clients. IT leaders looking for broader use-case adoption might end up with multiple products.
- ActiveScale is lacking in feature parity compared to its competitors. For example, it is missing features such as data deduplication and compression, QoS, mixed flash support, NFSv4 and SMB, hybrid cloud integration, and dual protocol access.

Qumulo

Qumulo is a Leader in this Magic Quadrant. Qumulo offers a portable software-defined file storage platform with data services that operates on-premises and in public clouds. The solution is designed for large-scale, high-throughput file workloads with built-in performance analytics and capacity management. Its primary market adoption has been in North America, where clients tend to be in the M&E, healthcare and public sectors. Over the past 12 months, Qumulo continued international expansion; hired several new executives in support of growth initiatives; launched a native, managed Azure service; offered support for AWS Outposts and introduced an all-NVMe appliance based on HPE servers. Qumulo is best suited for commercial HPC, analytics and hybrid cloud storage.

Strengths

- Qumulo's public cloud file service integration and hybrid cloud workflow enablement, as well as flexible consumption models on-premises and in public clouds, are well-aligned with the emerging needs of customers.
- Qumulo's number of customers with over one petabyte storage capacity under management rivals its larger competitors, underscoring client confidence in the use of Qumulo for multipetabyte requirements.
- Qumulo customers are expressing high satisfaction in ease of management, upgrades, and overall vendor support and responsiveness to their needs.

Cautions

- Compared to larger industry leaders in its peer group, Qumulo is a lesser-known brand name, requiring business validation prior to engaging in sales and proof-of-concept (POC) activities.
- Qumulo is dependent on its server and cloud OEM partners for global sales operations and marketing initiatives.

- Qumulo does not support nondisruptive removal or downsizing nodes per cluster, which could limit its value as an automated, managed file-as-a-service offering.

Red Hat

Red Hat is a Visionary in this Magic Quadrant. Red Hat Ceph Storage supports block, object and file workloads. Red Hat also sells a container storage product, OpenShift Data Foundation, which is based on Red Hat Ceph Storage. Vendor operations are focused in North America and Europe, and its clients tend to be large enterprises, telecom and financial services organizations. Over the past year, Ceph Storage has added an object lock WORM feature and FIPS 140-2 cryptography support, and has expanded key management integration. Recent Ceph Storage releases include tiering within a cluster by disk media type, improved performance for small object workloads and added CephFS file system support with NFS protocol-based access. Red Hat Ceph Storage is best suited for content delivery and cloud-native applications.

Strengths

- The IBM relationship has dramatically expanded Red Hat's access to the global market and traditional enterprises.
- Red Hat's proven record in open-source software-defined technologies and broad ecosystem support are attracting I&O leaders looking for flexible software-only storage powered by community-driven innovation.
- Red Hat Ceph Storage provides a versatile unified storage platform suitable for a variety of cloud-native application use cases.

Cautions

- As Red Hat positions Ceph Storage as part of its platform and data services, Gartner clients rarely shortlist it for consideration as a stand-alone enterprise storage platform for unstructured data.
- Red Hat Ceph Storage scores lower in manageability and support and services, compared to the competition.
- Red Hat's CephFS file system services are nascent and have low market adoption, limiting product applicability as a general-purpose file system.

Scality

Scality is a Leader in this Magic Quadrant. The Scality RING storage solution runs on-premises and extends to the public cloud. Scality offers integrated file and object storage for high-capacity unstructured data workloads and runs as software on commodity hardware. Its operations are focused in North America, EMEA and the Asia/Pacific region, and its clients tend to be distributed among almost every vertical. Over the past 12 months, Scality ported its SOFS file architecture to

Azure, announced ARTESCA as a separate lightweight container native object storage solution and added support for flash media. It also started offering several new features such as object locking, smart tiering and a new NFSv4 stack. Scality RING is best suited for multipetabyte geographical deployments of unstructured data for content distribution, media, backup and archiving requiring an SDS solution.

Strengths

- Scality's hybrid cloud integration and integrated capability to deliver file and object storage into a single solution remain attractive differentiators.
- Scality has a proven track record for very large-scale object and file-storage-based geographical distributed solutions across almost every vertical.
- Gartner clients are positive about the quality of sales, presales and technical support for Scality RING.

Cautions

- Scality's global installed base is relatively small compared to other Leaders in this Magic Quadrant, which can raise issues for long-term continuity.
- Scality RING is a lesser fit for smaller deployments, as it is best suited for large-scale multipetabyte geographical deployments.
- Gartner clients describe Scality as a "more complex" system that requires a higher level of expertise to deploy and manage.

Weka

Weka is a Visionary in this Magic Quadrant. WekaFS is software-defined distributed file system that can be deployed on-premises through OEM partners and in the public cloud, and has a parallel file system client. WekaFS was first released in 2017 and designed as an NVMe-based, I/O-intensive, low-latency distributed file system that can also extend to object storage in a single namespace on-premises or in the public cloud. Its operations are focused in North America and Europe. Over the past 12 months, Weka added Hitachi Vantara and Supermicro appliances, and expanded operations in the Asia/Pacific region. The latest product release includes RDMA networking, quota enforcement, S3 and autoscaling support in AWS. Weka is focusing on deployments for AI/machine learning (ML), financial analytics, life sciences and other HPC on-premises and in the public cloud.

Strengths

- Weka's software-only offering allows customers a choice of implementation on either commodity servers on-premises or at the edge through many server OEM vendors, or as SDS in the public cloud.

- Weka's multiprotocol capabilities and the ability to extend its file system to the object storage provides better scalability and cost-efficiency compared to traditional all-flash file systems, and enables hybrid cloud storage capabilities.
- Gartner clients are giving high grades to product performance as well as the overall simplicity of management compared to the competition.

Cautions

- As a relative newcomer to the unstructured data market, Weka still has limited brand awareness, traction and global reach, compared to the market leaders.
- Gartner clients do not often consider WekaFS as a replacement for a general-purpose file system or object storage, as it primarily targets specific performance-oriented use cases.
- The product is still missing some enterprise features (such as NFSv4, compression and deduplication support) and currently can only be deployed in the AWS public cloud.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Added

- Nutanix
- Weka

Dropped

- Caringo

Inclusion and Exclusion Criteria

To qualify for inclusion, vendors must meet *all* of the following requirements:

- The vendor must have recognized product revenue above \$10 million over the last four quarters (as of May 2021) for distributed file system and/or object storage solutions between 1 May 2020 and 30 April 2021 *and* should have at least 75 production customers each consuming more than 500TB raw capacity through either distributed file or object storage protocols only. Vendors must provide reference materials to support this criterion. This may require proof in the form of a

confidential list of representative customers outlining deployed capacity per product (75 customers with more than 500TB each). If customer names cannot be shared, then they can be anonymized as, for example, “large manufacturing company” or “small service provider.”

- The product must be installed in at least three (out of four) major geographies. Vendors must provide evidence of a minimum of 20 production customers brought to revenue in three out of four geographies (North America, EMEA, Asia/Pacific and South America). This may require proof in the form of a confidential list of representative customers from diverse geographies (20 customers of at least 500TB each in three out of four geographies). If customer names cannot be shared, then they can be anonymized.
- The product should be deployed across at least five out of the seven use cases that are outlined in Critical Capabilities for Distributed File Systems and Object Storage. Vendors must provide reference materials to support this criterion.
- The product must be designed for primarily on-premises workloads and not as a pass-through solution where data will be permanently stored elsewhere.
- Products should not be offered exclusively as a service.
- The vendor should own the storage software intellectual property and be a product developer. If a product is built on top of open-source software, the vendor must be one of the top 10 active contributors to the community (in terms of code contribution over the past 12 months).
- The vendor must have a product including features and capabilities generally available before 5 April 2021 that meet the criteria listed below under Packaging and Product Capabilities.

Packaging:

- The product must be sold as either an appliance or software-based storage solution.
- The product must be available for purchase and consumed as a stand-alone file- and/or object-storage-only product, and not as part of an integrated, converged or hyperconverged system with a compute and hypervisor bundle.

Product Capabilities:

- The product must have file and/or object access to the common namespace/file system.
- The product must have a fully distributed architecture where data and metadata are distributed, replicated or erasure-coded over the network across multiple nodes in the cluster. The product must have the ability to handle disk, enclosure or node failures in a graceful manner, without impacting availability.

- The product must be a single file system capable of expanding beyond 500TB.
- The product must have a global namespace capable of 2PB expansion.
- The cluster must span more than four nodes.
- The product must offer support for horizontal scaling of capacity and throughput in a cluster mode or in independent node additions with a global namespace/file system.

Note: A fully distributed architecture is a distributed computing architecture in which each node is independent and self-sufficient, and there is no single point of contention across the system. More specifically, none of the nodes share memory or disk storage. Distributed design systems are typically contrasted with systems that keep a large amount of centrally stored state information, whether in a database, an application or a metadata server, or any other similar single point of contention.

Honorable Mentions

Gartner tracks more than 20 vendors in this market. Fifteen vendors met the inclusion criteria for this Magic Quadrant, but the exclusion of a vendor does not mean that the vendor and its products lack viability. Following are several noteworthy vendors that did not meet all inclusion criteria, but that could be appropriate for clients, contingent on requirements.

Cohesity: Cohesity SmartFiles is a distributed file system product that provides an immutable file system with an integrated data and apps experience, and that acts as a multitier data manager. In addition to serving file and object data, end users can run antivirus, content search, file audit and analytics directly on the Cohesity platform.

MinIO: MinIO is an open-source-based distributed object storage suite. It is a software-only product and runs on industry-standard hardware in the public cloud. MinIO has many partnerships with storage and independent software vendor (ISV) providers, and is focused on performance-oriented workloads. It is fully integrated with Kubernetes, which makes it very popular among DevOps organizations.

VAST: VAST is a distributed file system and object appliance that is designed for large-scale data center deployments. Product data services and data storage nodes are connected via NVMe over fabric (NVMe-oF) protocols to enable higher scale, lower latency and global efficiency algorithms. New storage algorithms are leveraging storage-class memory to improve latency, resilience and throughput, and quad-level cell (QLC) flash media technology to improve flash economics.

Evaluation Criteria

Ability to Execute

We analyze the vendor’s capabilities across broad business functions. Ability to Execute reflects the market conditions and, to a large degree, it is our analysis and interpretation of what we hear from the market. Gartner analysts evaluate vendors on the quality and efficacy of the processes, systems, methods and procedures that enable IT provider performance to be competitive, efficient and effective, and to positively impact revenue, retention and reputation within Gartner’s view of the market.

Table 1: Ability to Execute Evaluation Criteria

Evaluation Criteria ↓	Weighting ↓
Product or Service	High
Overall Viability	High
Sales Execution/Pricing	Medium
Market Responsiveness/Record	High
Marketing Execution	Low
Customer Experience	High
Operations	Low

Source: Gartner (October 2021)

Completeness of Vision

Completeness of Vision distills a vendor’s view of the future, the direction of the market and the vendor’s role in shaping that market. We expect the vendor’s vision to be compatible with our view of the market’s evolution. A vendor’s vision of the evolution of the data center and the expanding role of distributed file and object storage are important criteria. In contrast with how we measure Ability to

Execute, the rating for Completeness of Vision is based on direct vendor interactions and our analysis of the vendor's view of the future.

Table 2: Completeness of Vision Evaluation Criteria

Evaluation Criteria ↓	Weighting ↓
Market Understanding	High
Marketing Strategy	Low
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	Medium
Innovation	High
Geographic Strategy	Medium

Source: Gartner (October 2021)

Quadrant Descriptions

Leaders

Vendors in the Leaders quadrant have the highest scores for their Ability to Execute and Completeness of Vision. A vendor in the Leaders quadrant has the market share, credibility, and marketing and sales capabilities needed to drive the acceptance of new technologies. Market leaders will typically be able to execute strongly across multiple geographies with products that cover both

distributed file systems and object storage offerings. They will also have consistent financial performance, broad platform support and flexible deployment models.

Challengers

Challengers are typically vendors with proven global presence and market achievement that only target a narrower subset of the market, or have not yet established themselves across the broader market for both the distributed file system and object storage areas. They have strong products as well as sufficient credible market position and resources to sustain continued growth in the future; however, they currently fall behind on influence and thought leadership for this market segment.

Visionaries

These are typically vendors that are focusing on strong innovation and product differentiation, but are smaller vendors with limited reach or achievement to date, or larger vendors with innovation programs that are still unproven. A vendor in the Visionaries quadrant delivers innovative products that address operationally or financially important end-user problems on a broad scale, but has not demonstrated the ability to capture market share or sustainable profitability.

Niche Players

Many distributed file system or object storage vendors will address a more narrow market niche, or they may be vendors with market programs that have not yet established their differentiation and/or execution ability. However, Niche Player vendors may address their specific market category and excel by focusing on specific market or vertical segments.

Context

This Magic Quadrant assesses vendors that sell products for unstructured data growth for enterprise data centers. The distributed file system and object storage market emerged as a response to the tremendous increase in unstructured data generation that is fueled by new business requirements. To address it, the storage platform has to be based on a scale-out software approach to enable seamless data growth with a strong emphasis on long-term data efficiency for cost optimization. I&O leaders seek distributed scale-out storage products to build new platforms based on software-defined approaches, where performance comes from hardware innovation of a commodity hardware layer and data resiliency comes from a scale-out software layer where data is distributed across multiple nodes.

Across many products in this market, vendors are providing appliances, software-only products and preintegrated storage systems to fit the needs of the different deployment strategies of enterprise end users. In addition, some of those products are enabling hybrid cloud workflows and being deployed in public cloud IaaS.

As the distributed file system and object storage market matures, storage software and hardware vendors are expanding their product portfolios to provide more differentiated and agile offerings.

New consumption models and procurement offerings are emerging to provide end users with different ways to purchase storage. Advances in software technology and commoditization of the hardware will make it possible for I&O leaders to enjoy web-scale economics and scalability of the storage platform for unstructured data growth.

Market Overview

The markets for distributed file systems and object storage have been merging. For that reason, Gartner publishes a single Magic Quadrant on the combined technology segments. The distinctions between the two segments are blurring, and buyers are already treating it as one market and requiring both file and object access for unstructured datasets. Moreover, the most innovative vendors now offer both object and file access to their unstructured data products.

IT leaders often decide between public cloud and on-premises infrastructure for given workloads. Organizational culture and sensitivity to security and governance mandates are typically the leading factors that enterprises consider when deciding whether to move applications and data to the public cloud or to keep them on-premises.

When customers choose to keep applications and data on-premises, they are increasingly choosing between object storage and file system products to accommodate the large sets of unstructured data. In many cases, customers seeking solutions in this market would be better-suited with a single product that has file and object personalities, so workloads can seamlessly interact with data using the most appropriate protocol for the specific task and environment:

- **Startups and Innovation:** Formerly risk-averse enterprises have become receptive to buying from storage startups that are using clean-sheet designs and a wealth of knowledge to build more efficient systems. This is illustrated by the popularity of products such as solid-state arrays, hyperconverged infrastructure, distributed file systems and object storage. Many large incumbent vendors are repositioning their distributed file systems for emerging AI workloads, while we are also seeing several emerging vendors like VAST and WekaIO specifically tackling the performance, scale and deployment flexibility improvements required for large-scale training and inference AI/ML workloads. Hyperconverged solutions like those from Nutanix and Cohesity are moving forward and replacing traditional network-attached storage (NAS) deployments, while also offering a single platform for multiple data service needs.
- **Choice in Deployment:** The vendors in the market for distributed file systems and object storage are offering mixed deployment options to give customers choices in how they deploy infrastructure. Common deployment options include turnkey high-density appliances or software-only options that can be deployed either on bare-metal industry-standard hardware as virtual machines or on Docker containers. Increasingly, vendors in this market are offering their products as SDS precertified to run on x86 industry-standard hardware.

- **Amazon S3 API Standardization:** The current object storage segment can be thought of as a two-sided market: There are providers of object storage protocols and consumers of these protocols consisting of applications. There were more providers than consumers until the Amazon S3 API became the de facto standard for object storage. Vendors deploying object storage platforms in enterprise data centers adopted Amazon S3, a protocol mainly used in the public cloud, because of the developer community that formed around it. Now there are many consumers and providers all using Amazon S3. The object storage market is finally in equilibrium. Interest in using public cloud services such as AWS has brought customer awareness to the object storage market. Software developers building Mode 2 web and mobile applications are sometimes asked to repatriate these applications back to enterprise data centers. Enterprise IT seeks control of applications and data, while software developers seek novel and efficient ways of programmatically interacting with infrastructure. The market for on-premises object storage products solves both of these.
- **Hybrid Cloud Storage:** The current unstructured storage market is evolving to embrace hybrid cloud workflows and capabilities as IT leaders are looking to take advantage of public cloud agility, efficiency and cloud computing capabilities. In 2021, Gartner saw more evidence of vendors not just adding tiering to the public cloud, but also enabling new use cases to leverage public cloud for rendering and analytics, and enabling application data bidirectional sharing between on-premises and public cloud locations.

Evidence

Placement on the Magic Quadrant for Distributed File Systems and Object Storage is based on Gartner's view of a vendor's performance against the criteria noted in this research. Gartner's view on vendor placement on the Magic Quadrant is heavily influenced by more than 1,200 inquiries and one-on-one meetings with Gartner clients regarding object storage and distributed file system solutions, conducted since the publication of the last iteration of this Magic Quadrant. Gartner also utilizes worldwide end-user surveys, Gartner conference session polling data, Gartner Research Circle polls and Gartner Peer Insights. The included vendors submitted comprehensive responses to Gartner's Magic Quadrant survey on this topic, which were used as the basis for subsequent vendor briefings and follow-up meetings, product demonstrations, and correspondence.

Additionally, this research drew input from other Gartner analysts, industry contacts and public sources, such as U.S. Securities and Exchange Commission filings, articles, speeches, published papers and public domain videos.

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered

natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service, and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.

**Learn how Gartner
can help you succeed**

Become a Client

© 2021 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)."

