

Magic Quadrant for Enterprise Storage Platforms

19 August 2026 - ID G00841627 - 31 min read

By Jeff Vogel, Joseph Unsworth, [and 2 more](#)

Enterprise storage platforms provide file, block and object data services for structured and unstructured workloads. Heads of infrastructure and IT operations can use this research to evaluate vendor platforms and support capabilities for modern storage infrastructure.

Strategic Planning Assumptions

- By 2029, consumption-based storage as a service (STaaS) will replace 50% of on-premises enterprise storage and data services infrastructure capital expenditure (capex), an increase from 25% in early 2026.
- By 2029, over 80% of on-premises unstructured data will be deployed on a consolidated storage platform instead of separate file and object products, up from 50% in 2025.
- By 2029, 70% of organizations will deploy purpose-built AI storage tiers specifically for large-scale AI inference and retrieval-augmented generation (RAG), up from less than 10% in 2026.
- By 2029, 100% of storage products will include cyberstorage capabilities focused on active defense beyond recovery from cyber events, up from 25% in early 2026.
- For 2027, heads of I&O must budget for 250% to 300% of their 2025 storage spend to accommodate surging memory and storage prices.

Market Definition/Description

Gartner defines enterprise storage platforms (ESPs) as the market consisting of products and services designed to unify support for diverse block, file and object storage workloads and use cases. ESP products and services include appliances, software-defined storage (SDS) and storage as a service (STaaS). An ESP includes data management and data storage services provided through a centrally managed, multidomain control plane. They enable organizations to leverage AI-powered telemetry for platform structured and unstructured workloads.

ESPs enable on-premises and hybrid IT operating models to increase flexibility, improve productivity and efficiency, improve asset management and enhance cyber resilience within the data storage environment. The foremost purpose of ESPs is to support structured and unstructured data storage workloads.

ESPs allow IT clients to consolidate multiple vendors across block, file and object workload infrastructure to simplify vendor management and reduce sourcing complexity. Use cases range from databases to high-performance file storage, and to backup and archiving. Global enterprise storage vendors have the technological, financial and operational means, platform strategies and roadmap initiatives to deliver ESP products and services.

ESP features leverage advanced AI-enabled telemetry that uses data collection and monitoring, in conjunction with machine learning technology, to enhance IT operations and support processes. The storage controller operating system is the foundation for platform innovation, providing a path to a software-defined storage (SDS) architecture that enables and simplifies hybrid IT platform services for multidomain infrastructure operations management.

Mandatory Features

- Support for structured and unstructured data workloads enabled by data access over Fibre Channel Protocol (FCP) for block and network file system (NFS), server message block (SMB) or Amazon Simple Storage Service (S3) protocols (protocol support must be native to the platform, not through a gateway or an external product)
- Primary block storage platform offerings to support structured data applications, including relational database management systems (RDBMSs), virtualization and container virtualization

- Block-storage platform with a minimum of 99.999% system availability, supported by active-active controllers and nondisruptive life cycle operations
- Either object and/or file-based data services to support unstructured data applications, including high-performance file and key-value repositories
- Native capabilities to protect against data loss, ensuring a minimum of 99.999% service-level assurance (SLA) data availability
- Support for a container management platform through a dedicated container storage interface (CSI) plug-in
- Either RAID or erasure coding data protection to withstand up to two simultaneous disk failures in the same fault tolerance group, redundancy and fault tolerance capabilities without any single point of failure
- API-centric central control plane for hybrid-cloud IT operations and storage management services, including orchestration and provisioning of distributed storage services

Optional Features

- Event intelligent solutions (EISs) (formerly artificial intelligence for IT operations [AIOps]) capabilities, typically provided by applying AI on telemetry data, leverage AI and machine learning to enable predictive insights for:
 - Prescriptive health management
 - Improved customer support
 - Proactive capacity and performance management
 - Nondisruptive workload simulation, placement and migration/tiering
- A distributed file architecture presented as a single namespace to scale the data stored across multiple servers/nodes
 - This feature supports the linear scaling of performance and capacity pooled across multiple storage nodes based on shared-nothing or shared-everything architecture principles. In addition, data and metadata should be distributed over multiple nodes in the cluster to handle availability and data protection in a self-healing manner.

- Multiple block and either file or object STaaS performance tiers that are available in a pay-for-use, consumption-based license or a pay-for-access subscription
- Support for a global file namespace that can span storage in multiple geographic locations at the edge, in on-premises data centers and on hyperscale public cloud platforms
- Cyberstorage capabilities, including advanced cyberthreat detection within the input/output data stream, real-time scanning and rapid recovery
- SDS deployment capability to operationalize native hybrid cloud and native public cloud storage infrastructure services with a minimum of one major public cloud
- A software-defined disaggregated architecture to support a multiprotocol storage architecture that logically separates the vendor's storage media hardware from the storage controller operating software managing the storage system and adjoining infrastructure
- A unified block and either file or object hardware architecture to provide an interchangeable and nondisruptive common workload infrastructure management environment
- Capacity arrays built on the highest-density quad-level cell (QLC) solid-state drive (SSD) flash media for price-sensitive applications or use cases:
 - Special purpose or next-generation QLC nonvolatile memory express (NVMe) SSD for enhanced performance, endurance, management or other critical storage services capabilities
 - Advantaged arrays using NAND and other components sourced directly from a semiconductor provider
- Data management services to create metadata classification enabling cost optimization, governance, mobility, analytics and security
- A nondisruptive data migration service that enables migration of array/SDS data from one platform to another with 100% data availability
- Data reduction capabilities with a guaranteed ratio for compression or deduplication
- Storage features that better support AI applications/workloads

Magic Quadrant

Figure 1: Magic Quadrant for Enterprise Storage Platforms



Magic Quadrant for Enterprise Storage Platforms



Gartner.

Vendor Strengths and Cautions

Dell Technologies

Dell Technologies is a Leader in this Magic Quadrant. PowerStore and PowerScale are its leading products for addressing structured and unstructured data workloads; both can be managed by using the Dell Automation Platform for monitoring and proactive remediation.

Dell's operations are geographically diversified and its clients tend to be midsize to large global enterprises. Dell is actively investing in private cloud infrastructure with PowerStore, which can be deployed as a private cloud foundation for Dell Private Cloud, and expanding its AI Data Platform with extreme-scale AI infrastructure, adding the purpose-built Lightning File System and deepening NVIDIA ecosystem integrations across the platform, while expanding its automated AI data preparation pipelines.

Strengths

- **Full-stack AI data platform foundation:** Dell's storage products form the foundation for its AI Data Platform, integrating directly with native data engines for AI orchestration, processing, analytics and search. This full-stack approach allows customers to manage the entire AI life cycle, from data ingestion and preparation to training and inference, within a single pipeline.
- **Comprehensive portfolio:** Dell's storage portfolio allows it to cater to any workloads — structured and unstructured, flash and hybrid, midrange to high-end, on-premises or public cloud — and to high-performance computing (HPC).
- **Global operations:** Dell's global supply chain, support and sales make it an attractive vendor for organizations with global operations, as well as regional enterprises of all sizes.

Cautions

- **Distinct architectures:** Dell's storage portfolio uses separate purpose-built platforms for structured and unstructured workloads, making fleetwide capacity optimization more complex, relative to that of leading vendors.
- **Architectural innovation gaps:** Dell trails market leaders in some architectural designs, specifically in native custom silicon development, hybrid cloud ecosystem and embedded, in-place AI microservices.
- **Deployment friction:** Some customers may encounter greater initial deployment complexity, depending on product configuration or environment. Day-to-day management can be challenged by inconsistent technical support resolution.

Everpure

Everpure, formerly Pure Storage, is a Leader in this Magic Quadrant. Its FlashArray and FlashBlade products are broadly focused on mission-critical databases, unstructured

enterprise AI and hybrid cloud storage. Its operations are geographically diversified, although it has limited exposure to emerging markets, and its clients vary in size from SMEs to large global enterprises, and encompass the financial services, healthcare and technology sectors. Everpure continues to invest heavily in its Enterprise Data Cloud vision, developing agentic autonomous management through its Fusion platform, advanced data intelligence following its acquisition of 1touch, and purpose-built hardware.

Strengths

- **Operations:** Everpure's operational acumen and long-standing flash technology partnerships underpin its all-flash array cost competitiveness and product availability assurances for customers. This strategic stability can allow IT decision makers to derisk their infrastructure rollouts with more predictable pricing and deployment schedules.
- **Exceptional customer experience:** Everpure's Pure1 AI-driven telemetry reduces administrative troubleshooting. It delivers extensive fleet management capabilities and leading proactive support resolution, which eases administrative burden for users.
- **Predictable consumption economics:** The Evergreen//One subscription model eliminates disruptive forklift upgrades, simplifying administrative overhead, and shields end users from supply chain volatility by delivering nondisruptive service-level assurance (SLA)-backed outcomes.

Cautions

- **Premium acquisition costs:** Cost-conscious enterprises may view Everpure as a high-cost option, particularly for capacity-centric use cases where competitors offer lower-cost HDD media and more economical support contracts. This can impede organizations prioritizing budget optimization over premium platform capabilities.
- **Platform-specific file namespaces:** Everpure lacks support for a cross-platform global file namespace. This requires end users to adopt the vendor's specialized, high-end appliance for performance scale.
- **Geographic presence:** Everpure lags behind competitors in customer diversification in emerging countries, making it difficult for customers to justify selection based on brand visibility, compared to established vendors in those regions.

Hitachi Vantara is a Visionary in this Magic Quadrant. Its Virtual Storage Platform (VSP) One portfolio targets mission-critical workloads, online transaction processing (OLTP), virtualization and data lakehouse environments. Its operations are geographically diversified and have a global customer base, primarily among large enterprises in regulated industries such as financial services, healthcare, operational technology (OT) and critical infrastructure. Hitachi Vantara's roadmap focuses on investing in VSP One to support high-performance GPU-driven workloads, expanding its quad-level cell (QLC) flash storage capabilities, and enhancing its VSP 360 management control plane with agentic-AI-powered autonomous management and operations.

Strengths

- **Mission-critical availability:** The VSP One platform features an active-active clustering architecture that ensures continuous operations and zero data loss for enterprise workloads. This resilient design is backed by a contractual 100% data availability guarantee.
- **Platform vision:** VSP 360 provides a single control plane, including AI-powered autonomous capabilities, to manage block, file and object storage. This unified cloud management approach simplifies administration, automates troubleshooting and provides fleet-wide observability for large deployments.
- **Strong global footprint:** Hitachi Vantara's worldwide operations enable reliable, regionally aligned support and services, backed by local technical expertise and a well-established global partner ecosystem.

Cautions

- **Lower market visibility:** Due to Hitachi Vantara's relatively lower market visibility, customers encounter a smaller pool of professionals with active training on the platform, compared to market leaders. Consequently, access to peer-to-peer community support and available talent may be more limited, which results in a greater reliance on the vendor for specialized assistance.
- **High upfront costs:** The initial hardware investment for Hitachi Vantara's midtier storage arrays is frequently cited by customers as expensive compared to comparable market products.
- **Limited AI storage awareness:** Hitachi Vantara struggles to gain wide adoption in AI training and inferencing environments because it lags market leaders in the breadth and depth of its dedicated AI data services, such as vector database integration.

HPE

Hewlett Packard Enterprise (HPE) is a Leader in this Magic Quadrant. Its HPE Alletra Storage MP B10000 and X10000 products enable the modular storage platform to deliver unified block, file and object storage for mission-critical workloads and consolidated unstructured data infrastructure, including AI workloads. HPE operations are geographically diversified and serve large global clients in the banking, technology and healthcare sectors. Its roadmap focuses on GreenLake Intelligence for autonomous agentic-AI operations, expanding X10000 scalability to address all needs of unstructured data, and delivering Data Fabric Appliances for integrated data intelligence and hybrid cloud data services.

Strengths

- **Enterprise storage platform:** HPE Alletra Storage MP architecture offers a fully disaggregated, modular scale-out storage platform supporting block, file and object workloads from a single console. This unified foundation allows performance and capacity to scale independently, removing legacy storage silos and simplifying planning and provisioning.
- **Autonomous AI operations:** HPE GreenLake integrates a multiagent AI architecture trained on global telemetry to deliver predictive remediation and natural language management. This creates an ecosystem capable of self-managing that eliminates reactive troubleshooting to uphold the platform's 100% data availability guarantee.
- **Consolidated data services:** HPE Alletra Storage MP serves as the common storage architecture across all HPE Data Services, underpinning both HPE Morpheus VM Essentials software and HPE Private Cloud AI. HPE GreenLake unifies compute, storage and hypervisors into a single control plane to streamline enterprise data workloads.

Cautions

- **GreenLake perception:** The perception that purchasing HPE storage mandates a full transition to the GreenLake hybrid cloud ecosystem can cause buyers to bypass HPE's independent storage products during vendor selection.
- **Advanced unstructured data services:** As a newer data platform, the X10000 currently lacks several mature features found in established offerings, such as automated failover, multisite active-active georeplication, and advanced NFS v4.1 capabilities.

- **Narrow hybrid cloud market traction:** Despite the availability of Alletra Block Storage in Azure and AWS, HPE struggles to gain hybrid cloud traction. Customer usage is largely confined to data protection, rather than diverse cloud-native applications.

Huawei

Huawei is a Leader in this Magic Quadrant. Its OceanStor Dorado All-Flash Storage, OceanStor Pacific Storage, and AI Data Platform products broadly focus on mission-critical, OLTP, high-performance unstructured data, and enterprise AI workloads. Its operations are principally in Asia/Pacific and across EMEA and Latin America, mostly serving large enterprise clients in the banking, telecommunications, government and healthcare sectors. Huawei continues to invest heavily in AI-native storage architectures featuring UnifiedBus networks, Context Memory Storage, and vertically integrated, high-density SSD hardware to mitigate supply chain volatility.

Strengths

- **High-density infrastructure design:** Huawei's deep vertical integration, custom-engineered packaging and industry-leading capacity density enable organizations to maximize compute and storage within a smaller footprint. This reduces data center space requirements, lowers power consumption and improves overall operational efficiency.
- **Future-ready cyber resilience:** Huawei has proactively integrated cryptographic controls to help organizations safeguard critical data from emerging threats, including proactive protection against future quantum-enabled decryption risks. This ensures sensitive data remains secure at rest.
- **Unified AI data platform:** Huawei's integrated AI data platform simplifies enterprise AI deployments by accelerating inference performance, maximizing GPU utilization and reducing cost per token. This helps organizations scale AI initiatives while improving return on infrastructure investments.

Cautions

- **Support automation gaps:** Huawei resolves fewer support issues autonomously with AI-driven telemetry than leading competitors. Some customers report occasional communication challenges with outsourced Tier 1 support teams.
- **Western cloud gaps:** Huawei's software-defined storage (OceanSDS) and cloud presence in Western hyperscalers lag market leaders.

- **Geopolitical supply constraints:** Geopolitical instability and export controls continue to constrain Huawei's operations outside of its home market. This creates challenging barriers to client adoption, support and hardware lock-in, in addition to limited flexibility to integrate best-of-breed solutions.

IBM

IBM is a Leader in this Magic Quadrant. Its portfolio, increasingly integrated under the IBM Fusion platform, includes IBM FlashSystem for block storage, and IBM Storage Scale and IBM Storage Ceph for unstructured data. It is broadly focused on enterprise AI, HPC and mission-critical workloads. It operates globally, primarily targeting regulated enterprises in financial services, government, healthcare and technology. Its future investments focus on agentic storage management through FlashSystem.ai, high-density arrays powered by FlashCore Module 5 (FCM5), and Content Aware Storage (CAS) to natively vectorize data for AI pipelines.

Strengths

- **Agentic storage management:** IBM integrates agent administrators, agentic APIs and conversational AI capabilities directly into its platform data path to manage, diagnose and remediate infrastructure issues. This allows administrators to orchestrate complex operations and provision workloads using natural language, significantly reducing manual overhead.
- **Hardware-accelerated cyber resilience:** IBM embeds AI-driven ransomware detection natively into its proprietary flash modules, delivering threat identification in under 60 seconds. Chip-level hardware integration provides verifiable and immutable recovery capabilities specifically tailored for highly regulated enterprises.
- **AI-ready unstructured data:** The content-aware platform capability securely provisions vectorized file and object data directly to enterprise AI inference pipelines, reducing the time and cost required to operationalize enterprise AI.

Cautions

- **Platform integration complexity:** IBM currently lacks a native unified platform, as a result of integration of distinct and separate purpose-built structured and unstructured products. This adds operational cost and complexity for users.

- **Administrative learning curve:** IBM's architectures can be burdensome for smaller organizations, requiring a steep learning curve for specific tasks such as system integration, performance tuning and ongoing upgrades.
- **Reactive customer support:** IBM's support services can result in prolonged back-and-forth communications when customers seek a resolution.

IEIT SYSTEMS

IEIT SYSTEMS is a Niche Player in this Magic Quadrant. Its AS/HF Series, AS13000G7 Distributed Storage System Series, AS3000G7 Storage Platform, and A9000 Storage Platform broadly focus on large-scale AI training and inference, enterprise digital transformation and smart healthcare. Its operations are mostly focused in China, Southeast Asia, South Korea and Eastern Europe, serving mostly large enterprises, intelligent computing centers and research institutions. IEIT SYSTEMS continues to invest heavily in AI scenario use-case-based storage-compute synergy and next-generation storage architectures, while expanding its localized delivery and manufacturing infrastructure across Asia, Europe and Latin America.

Strengths

- **Cost-optimized infrastructure:** IEIT SYSTEMS' strong hardware R&D and supply chain capabilities allow it to deliver highly cost-effective capacity scaling for large enterprise datasets without premium markups.
- **Accelerated AI pipelines:** The IEIT SYSTEMS platform natively integrates Nvidia GPUDirect Storage (GDS) with external KV Cache offloading. This effectively uses storage to substitute for expensive GPU compute to prevent GPU bottlenecks and reduce time to first token (TTFT) for large-scale AI inference.
- **Turnkey storage-compute delivery:** IEIT SYSTEMS provides an integrated delivery model combining AI computing servers, parallel file storage and object storage, with joint tuning specifically optimized for large models. This simplifies multivendor troubleshooting and accelerates AI deployment.

Cautions

- **Concentrated market presence:** IEIT SYSTEMS' brand recognition and primary operations remain heavily concentrated within its domestic Chinese market and select Asia/Pacific regions. This can create possible support issues for globally distributed enterprise deployments.

- **Opaque product roadmaps:** Limited insights into IEIT SYSTEMS' product roadmap create uncertainty regarding its commitment to the platform's long-term architectural evolution and its ability to seamlessly support future workload requirements.
- **Reactive, manual support:** IEIT SYSTEMS' customer support model is highly manual compared to market leaders, resolving few issues autonomously via telemetry. Additionally, IEIT SYSTEMS relies on its R&D engineers for manual performance tuning.

NetApp

NetApp is a Leader in this Magic Quadrant. Its ONTAP and StorageGRID products are broadly focused on providing a unified data foundation for block, file and object workloads across hybrid multicloud environments. Its operations are geographically diversified, and its clients tend to be large global enterprises and cloud-first organizations. NetApp continues to actively invest in disaggregated storage architectures and its AI Data Engine, with the goal of bringing vectorization and intelligent data preparation microservices directly to the storage layer.

Strengths

- **Cloud integration:** NetApp minimized architectural fragmentation by embedding its storage operating system natively within AWS, Azure and Google Cloud. This helps buyers manage hybrid data seamlessly while drawing down their existing cloud financial commitments.
- **In-place AI:** NetApp's platform runs containerized microservices directly on the array to execute metadata extraction and vectorization. This allows enterprises to prepare unstructured data for AI workflows without the cost and risk of moving massive datasets to external compute silos.
- **Supply resilience:** By manufacturing in Mexico under USMCA exemptions and diversifying its component supply chain away from China, NetApp insulates enterprise buyers from tariff volatility and offers more predictable hardware delivery timelines than its competitors.

Cautions

- **High total cost:** NetApp's premium pricing model, expensive ongoing maintenance contracts and high capital expense regarding expanded physical capacity make the platform difficult to justify for premium enterprise configurations.

- **Object storage fragmentation:** ONTAP S3 translates object APIs onto a traditional file system, limiting scale and multiregion distribution. To get hyperscale, geodistributed object repositories, buyers must purchase NetApp's separate StorageGRID product, which introduces architectural and technological complexity.
- **Management interface complexity:** NetApp's graphical user interface can be confusing and challenging to navigate for advanced workflows, requiring administrators to face a steep learning curve, often relying on command-line interface inputs to fine-tune the system.

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.

Dropped

DDN

Inclusion and Exclusion Criteria

Magic Quadrant Inclusion Criteria

To qualify for inclusion, providers must meet all of the following criteria and provide evidence, if requested, to support some or all of the claims henceforth:

- Vendors must offer block and either file or object storage services, each with an API-centric central control plane with integrated AI telemetry, capable of provisioning, orchestration and platform life cycle management in a nondisruptive operations manner.
- The services must contain a standardized common API catalog with the broader ecosystem of programmatic capabilities to interact with hybrid infrastructure and services to provide visibility and control for consistent data management and orchestration.

- Vendors must have a minimum total product and services revenue of \$325 million (excluding support and maintenance), with at least 25% of total revenue coming from unstructured data product revenue. In addition, unstructured data product revenue must have a minimum year-over-year growth of 15%, as reported through 31 May 2026.
- Vendors must have a minimum of 500 active production block customers and a minimum of 200 active production file or object storage customers, with these services being generally available (GA) from 1 January 2026. New features and capabilities must be generally available by 31 May 2026.
 - Each active customer must have a minimum of 100TB of raw block storage and 500TB of raw file or object storage.
- Vendors must support active customers in three of five major geography regions (e.g., North America; Europe, the Middle East and Africa; Asia/Pacific; Japan; and Latin America). In addition, vendors must include at least 25 customers of block and either file or object storage services for each major geography region supported.
- Vendors must support a minimum of five out of seven use cases, with mandatory support for online transaction processing (OLTP), container management and virtualization platforms. Vendors must show evidence for customers deployed across at least five out of seven use cases. The use cases are as follows:
 - Online transaction processing
 - Virtualization and container management platforms
 - High-performance file
 - Object-native applications
 - Hybrid platform services
 - Hybrid cloud storage
 - Artificial intelligence
- Vendors must own (not manufacture) and develop the core storage operating system, including the block store and file or object system.

- Vendors must offer a minimum of two integrated, managed and supported data storage services — either organic or partner-branded — for backup or disaster recovery (DR), cybersecurity, archive or database-related services.
- Vendors must offer data management capabilities between on-premises and at least one of these top two public clouds by global revenue: AWS or Azure.
- Vendors must offer block and either file or object as a storage service (STaaS).

Magic Quadrant Exclusion Criteria

- ESP primary storage SDS (software-defined storage) options are excluded from this market if SDS is part of a hyperconverged infrastructure (HCI) solution or if the storage software is not available independent of hardware for on-premises usage.
- Public cloud storage services offered by infrastructure as a service (IaaS) or platform as a service (PaaS) providers are excluded, including those provided by cloud service providers (CSPs) that integrate and offer the vendor's core product intellectual property.

Honorable Mentions

The providers that are most relevant to our clients were selected for evaluation in this Magic Quadrant. However, the decision not to evaluate a provider does not mean that the provider lacks viability. The following are noteworthy providers not included in the formal analysis.

These providers could be appropriate for clients, contingent on their requirements:

Infinidat, a Lenovo Company: Infinidat earns an honorable mention in the Magic Quadrant for Enterprise Storage due to its strong reputation among large enterprises, combining high-performance, cyber-resilient storage with a differentiated architecture focused on availability and efficiency. Lenovo completed the acquisition of Infinidat on 9 April 2026, after receiving the required regulatory approvals, with Infinidat part of Lenovo's Infrastructure Solutions Group. Infinidat continues to stand out through exceptional customer satisfaction and recognition in enterprise storage evaluations, making it an alternative to the larger incumbent vendors.

Evaluation Criteria

Ability to Execute

Ability to Execute reflects our assessment of a vendor’s capability to deliver its vision, meet customer expectations, and achieve sustained success in the enterprise storage platforms market. In evaluating vendor organizations and their platform offerings, we consider a range of factors that demonstrate operational effectiveness, market performance, customer outcomes and overall competitiveness. The table below identifies the specific criteria included in our assessment and the weighting assigned to each, reflecting our perspective on their relative importance within the current and evolving market landscape.

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

Source: Gartner (August 2026)

Completeness of Vision

Completeness of Vision reflects our assessment of a vendor’s ability to understand market dynamics, anticipate future customer needs, and translate that understanding into a clear and differentiated strategic direction. In evaluating vendor organizations and their platform offerings, we consider factors such as market insight, innovation, product and business

strategy, and the degree to which the vendor is positioned to influence and adapt to market evolution. The table below identifies the specific criteria included in our assessment and the weighting assigned to each, reflecting our perspective on their relative importance to long-term success in the enterprise storage platforms market.

Completeness of Vision Evaluation Criteria	
<i>Evaluation Criteria</i>	<i>Weighting</i>
Market Understanding	High
Marketing Strategy	High
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	Medium
Innovation	High
Geographic Strategy	Medium

Source: Gartner (August 2026)

Quadrant Descriptions

Leaders

Vendors in the Leaders quadrant have the highest composite scores for their Ability to Execute and Completeness of Vision. A Leader has broad market share across major geographies, exceptional brand awareness, recognized outstanding support, robust financial performance and established enterprise credibility. These vendors are market-driven — they

anticipate customer demands over time and make the long-term capital and R&D investments needed to drive platform-native outcomes, advanced technologies, capabilities and extensive ecosystem partnerships. Leaders demonstrate a deep, proactive understanding of enterprise storage platform (ESP) market needs, showing how product features, combined with platform business model innovations, transform and modernize IT operations. They are recognized innovators and thought leaders, with well-articulated roadmaps that global enterprises can depend on for their long-term hybrid infrastructure strategies.

Challengers

Vendors in the Challengers quadrant participate in the broad enterprise storage platform market and execute well enough on traditional high-performance metrics — such as physical throughput, raw IOPS, and established controller-level reliability — to represent a serious threat to vendors in the Leaders quadrant. Challengers have enterprise-class storage product features with which to compete, but lag behind Leaders in platform business model or technological innovation. Challengers have not yet demonstrated the long-term and consistent investments in leading-edge, platform-enabled product capabilities, or the business model, marketing, and sales strategies needed to transition from selling discrete, siloed storage boxes to delivering policy-governed, STaaS-driven data outcomes.

Visionaries

Vendors in the Visionaries quadrant are providing leading market-driven infrastructure platform features and innovative platform capabilities across the spectrum of the hybrid IT operations platform stack offerings to address operationally or financially important end-user demands. Visionaries are recognized by enterprise storage clients and peers for their ability to continuously demonstrate compelling thought leadership and an innovative view of the future for modern IT operations, often introducing disruptive platform initiatives, supporting technologies or approaches. Visionaries have not yet demonstrated the ability to capture meaningful platform market share through broad market-driven sales and marketing initiatives that underpin concrete commercial traction.

Niche Players

Vendors in the Niche Players quadrant are often narrowly focused on specific geographical markets, vertical industry segments or limited use cases. This quadrant may also include vendors that are ramping up their product and platform offerings, or larger vendors that are

having difficulty developing and executing on their vision against new platform market demands. Their overall execution remains constrained by lower global brand visibility outside of their brand reputation.

Context

As enterprise storage solutions scale, client organizations are actively shifting away from traditional IT storage budgeting and fragmented storage hardware sourcing to platform-native data services consumption. This Magic Quadrant for Enterprise Storage Platforms places its primary emphasis on an ESP's ability to serve as an active data control plane that natively unifies block, file and object workloads under a centrally managed, hybrid-cloud operational control plane.

To navigate this period of intense industry transformation, heads of I&O must execute the following four strategic actions:

- **Shift sourcing from capex to SLA-backed STaaS:** Mitigate hardware cost inflation and technology obsolescence by transitioning to consumption-based STaaS models. Ensure these agreements are contractually assured by enforceable SLA outcomes that are managed by platform vendors.
- **Eliminate ETL complexity with in-place AI architectures:** When designing infrastructure for enterprise AI and RAG environments, avoid the high egress fees, transport latency, and compliance risks of manual extract-transform-load (ETL) data pipelines. Prioritize disaggregated scale-out platforms that bring AI workloads directly to the data by executing on-array metadata indexing, semantic tagging and vector generation.
- **Deploy hypervisor-flexible storage to reduce platform lock-in:** To systematically manage and mitigate hypervisor restructuring risks, avoid storage platforms that tightly bind data to a single hypervisor engine. Prioritize platforms that deliver hypervisor-agnostic deployment flexibility and native VM-to-container migration tools (supporting Red Hat OpenShift, Nutanix AHV, Microsoft Azure Local, and KubeVirt) to enable nondisruptive, in-place workload transitions.
- **Mandate active cyberstorage as a primary security requirement:** Avoid passive postevent recovery models that only scan snapshots hours after an attack has occurred. Instead, prioritize storage platforms that integrate native, in-line anomaly detection directly within

the storage data path. Select platforms that autonomously trigger immutable snapshot locks to isolate compromised hosts at the first sign of anomalous encryption behavior.

Market Overview

ESP Market Undergoes Unprecedented Transformation

The enterprise storage platforms market is undergoing a structural transformation unlike any before, driven by four major market forces: extreme supply chain volatility and costs, the transition to production-scale artificial intelligence, hypervisor replatforming, and cyber-data-resilient architectures.

Extreme Supply Chain Volatility

The rapid build-out of enterprise AI factories has triggered a massive imbalance in supply and demand, resulting in severe shortages and drastic price increases for critical silicon and memory components. Over the past 12 months, media costs have risen sharply, resulting in the price per gigabyte of enterprise SSDs increasing by more than 15 times that of nearline HDDs. A storage-constrained environment, with elevated costs, is projected to persist into at least 2028, fundamentally reshaping how I&O leaders budget, procure and manage storage.

The extreme cost inflation has elevated the importance of HDDs within tiered and hybrid architectures, ensuring that HDD and even tape will remain essential through 2030. While nearline HDDs are also experiencing severely extended lead times, with some high-capacity models reported as being backordered for a year or more, enterprises are actively reprioritizing tiered designs. Under this model, high-density, performance-optimized QLC solid-state media serve as the primary ingestion tier for active data lakes, while HDDs perform as high-capacity, cost-effective warm storage rather than deep-cold archives. Intelligent analytic platforms that intimately understand, predict and automate data placement from core to cloud to edge will be fundamental for I&O leader cost optimization efforts.

Production-Scale Artificial Intelligence

The primary performance bottleneck in the accelerated computing era has shifted from GPU compute to the storage data layer, changing buyer evaluation metrics from traditional raw capacity to active, multidimensional performance indexes like cost per token and tokens per watt hour. To prevent expensive GPU clusters from sitting idle due to data readiness delays,

modern enterprise storage is transitioning from isolated hardware arrays into intelligent, unified data fabrics that establish a federated global namespace spanning core, edge, and hybrid-cloud data environments. These fabrics allow organizations to run secure, in-place AI and RAG pipelines directly where the data resides by embedding native data intelligence, metadata cataloging, semantic tagging and vectorization directly within the primary storage path.

Hypervisor Replatforming

Workload replatforming and hypervisor modernization due to widespread licensing restructuring and pricing changes are actively disrupting the virtualization market, with Gartner projections indicating that a meaningful percentage of enterprise VMware workloads will migrate to other platforms by 2028. This transition has opened a major displacement window, forcing IT leaders to evaluate virtualization and storage platforms as a combined entity. Consequently, buyers are heavily prioritizing storage platforms that offer multihypervisor flexibility, native container interoperability and automated VM-conversion tools to eliminate data migration friction.

Cyber Data Resilience

The convergence of storage and cyber data resilience has fundamentally changed the role of enterprise data storage in the era of heightened cyberthreats. The boundary between network security and data resilience has dissolved, shifting the ownership of disaster recovery directly to CISOs and risk leaders. Organizations must be able to defend against highly automated, AI-generated ransomware attacks that can encrypt entire file systems in minutes. Storage is now an integral line of defense. This has catalyzed a major market transition toward active cyberstorage capabilities. It requires real-time in-line encryption detection, secure isolated recovery environments (IREs), and postquantum cryptography (PQC) natively on the management plane, and replication links to protect sensitive data against future decryption risks.

⊕ Evidence

⊕ Evaluation Criteria Definitions

Unlock more AI value

With unmatched, proprietary AI insights from Gartner

[Explore the AI Hub ↗](#)

© 2026 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](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

[About Us](#) [Search Careers](#) [Newsroom](#) [Policies](#) [Site Index](#) [Featured Insights](#)
[Contact Us](#) [Send Feedback](#)

Gartner

© 2026 Gartner, Inc. and/or its Affiliates. All Rights Reserved.