GIGAOM

In partnership with:

Hitachi Vantara

Radar for Object Storage v7

From Archive to AI Factory: How Object Storage Became the Foundation of Modern Infrastructure | Whit Walters

CLOUD, INFRASTRUCTURE & MANAGEMENT

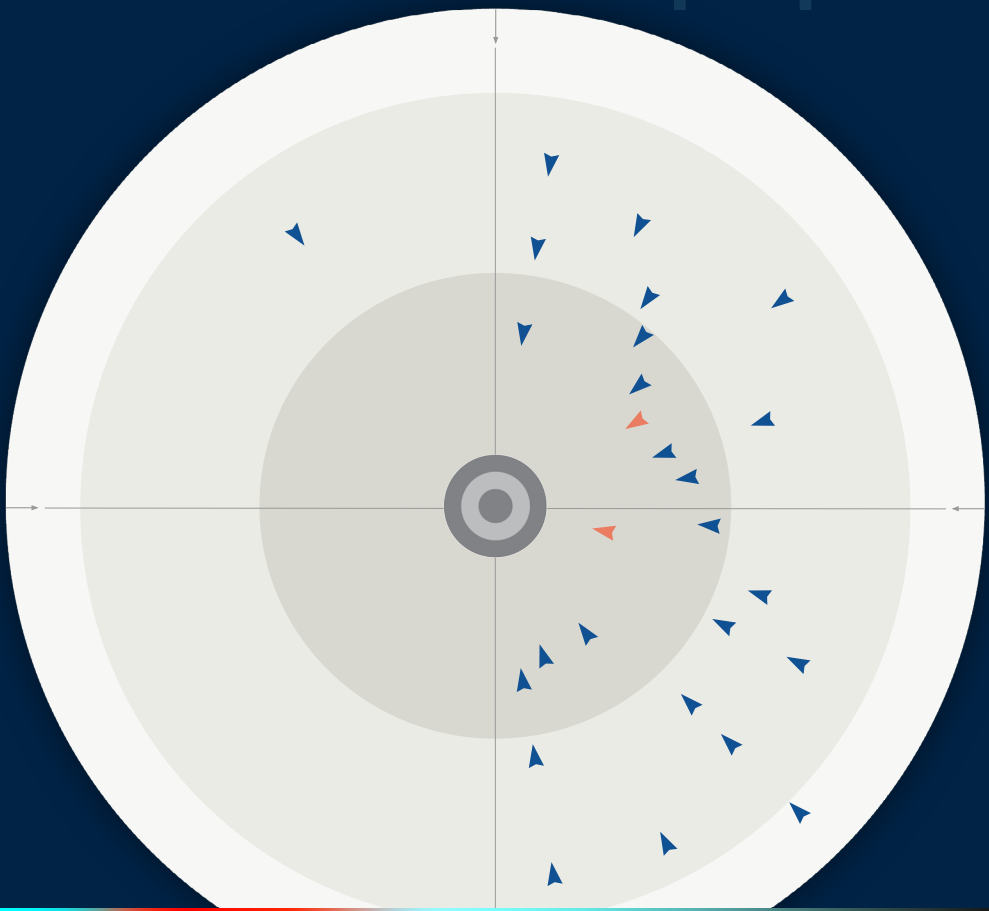


GIGA
OM

Radar for Object Storage v7

Whit Walters

Table of Contents

01 Executive Summary

This GigaOm Radar report is commissioned by Hitachi Vantara

OBJECT STORAGE HAS EVOLVED fundamentally from its origins as a "cheap and deep" repository for archival data into a primary storage tier capable of supporting high-performance workloads. It is now the dominant architectural standard for unstructured data, serving as the backbone for modern application stacks, cloud-native development, and massive-scale analytics. This technology provides a flat, scalable namespace that allows organizations to store petabytes to exabytes of data with rich metadata, ensuring accessibility and durability across on-prem, edge, and hybrid cloud environments.

For CxOs and IT leaders, the business imperative for adopting modern object storage is driven by the need to monetize and protect data assets. As organizations aggressively adopt AI and ML, storage must deliver not just capacity but the high throughput and low latency required to saturate GPUs and accelerate training cycles. Furthermore, with the rising threat of cyberattacks, object storage now serves as a critical layer of defense.

Features like immutability (Object Lock) are essential for regulatory compliance and ransomware resilience.

This report focuses on enterprise object storage solutions that can be sold and deployed as standalone offerings, whether as software-defined storage (SDS), a turnkey appliance, or a dedicated storage-as-a-service platform. Solutions must support the Amazon S3 API as a de facto standard and demonstrate capabilities suitable for enterprise production environments.

This is our seventh year evaluating the object storage space. This report builds on our previous analysis and considers how the market has evolved over the last year.

This GigaOm Radar report examines 27 of the top object storage solutions and compares offerings against capabilities (table stakes, key features, and emerging features) and nonfunctional requirements (business criteria). It provides an overview of the market, identifies leading object storage offerings, and helps decision-makers evaluate these solutions so they can make a more informed investment decision.

02 Market Categories and Deployment Types

TO HELP PROSPECTIVE CUSTOMERS find the best fit for their use case and business requirements, we assess how well object storage solutions are designed to serve specific target markets and deployment models (**Table 1**).

For this report, we recognize the following market segments:

- **Small-to-medium business (SMB):** Solutions targeting this segment prioritize simplicity, ease of deployment, and low total cost of ownership (TCO). These organizations typically have limited IT staff and require "set it and forget it" operations, often valuing turnkey appliances or simple cloud services over complex, highly tunable architectures.
- **Large enterprise:** This segment demands massive scalability, granular security— role-based access control (RBAC), MFA, Object Lock—and deep integration with existing identity management and analytics ecosystems. Buyers in this space prioritize high availability, multisite resilience, and the ability to handle diverse workloads, from archival to high-performance AI training, within a single namespace.
- **Cloud service provider (CSP):** This segment includes organizations building large-scale public or private cloud infrastructure to sell storage services directly to end users. Key requirements include massive, linear scalability to exabytes, fully automated API-driven provisioning (infrastructure-as-code), and advanced multitenancy with strict quality-of-service (QoS) controls to prevent "noisy neighbor" issues. CSPs prioritize white labeling capabilities and granular billing integration to monetize storage capacity seamlessly.
- **Managed service provider (MSP):** MSPs manage IT services for multiple distinct clients, often bundling storage with other offerings like backup-as-a-service (BaaS) or disaster recovery (DR). Buyers in this segment look for multitenant management consoles that simplify monitoring across a distributed client base. They value ease of deployment, robust reporting for chargeback/showback, and integration with remote monitoring and management (RMM) or professional services automation (PSA) tools over raw hyperscale architecture.

In addition, we recognize the following deployment models:

- **Physical appliance:** The vendor provides a turnkey solution with integrated hardware and software. This model simplifies deployment and support by offering a "single throat to choke" for the entire stack. It is often preferred for on-prem data centers where specific performance tuning (such as via all-flash nodes) or high-density capacity is required without the complexity of component compatibility testing.
- **Virtual appliance:** The solution is delivered as a preconfigured VM image designed to run on standard hypervisors (VMware, KVM, Hyper-V). This allows organizations to leverage existing virtualization infrastructure for ease of management and rapid provisioning, making it ideal for edge locations or branch offices where dedicated hardware is unnecessary.

- **Public cloud image:** The solution is available as a prepackaged machine image (such as an Amazon Machine Image [AMI], Virtual Hard Disk [VHD], and others) in public cloud marketplaces, including AWS, Azure, and Google Cloud. This allows organizations to deploy the vendor's specific object storage software stack on top of public cloud infrastructure, ensuring feature parity (such as consistent encryption or replication policies) between on-prem and cloud environments.
- **Software only:** The vendor provides the object storage software as an installable binary or container that runs on the customer's choice of commodity bare metal servers and operating systems. This model offers the highest flexibility and hardware independence, allowing IT architects to optimize the underlying infrastructure for specific cost or performance metrics (like reuse of existing servers).
- **Software-as-a-service (SaaS):** The solution is consumed as a fully managed service by which the vendor handles all infrastructure, patching, and scaling. Customers simply consume the API endpoint and pay based on usage (capacity, egress, and so on). This model eliminates operational overhead and accelerates time to value, appealing to organizations that prefer an OpEx model over managing infrastructure.
- **Self managed:** The customer deploys and operates the object storage software within its own infrastructure (on-prem, colocated, or in a private cloud), retaining full administrative control over configuration, upgrades, and capacity planning. This model appeals to organizations with strict data sovereignty requirements or existing operational teams that prefer direct control over the storage lifecycle.

Table 1. Vendor Positioning: Target Market and Deployment Model

	TARGET MARKET				DEPLOYMENT MODEL					
PLATFORM	SMB	LARGE ENTERPRISE	CSP	MSP	PHYSICAL APPLIANCE	VIRTUAL APPLIANCE	PUBLIC CLOUD IMAGE	SOFTWARE ONLY	SAAS	SELF MANAGED
Cloudian	-	✓	✓	✓	✓	✓	✓	✓	-	✓
Cohesity	-	✓	✓	✓	✓	✓	✓	✓	✓	✓
Cubbit	✓	✓	✓	✓	-	-	-	✓	✓	✓
DataCore Software	✓	✓	✓	✓	✓	✓	✓	✓	-	✓
DDN	-	✓	✓	✓	✓	✓	✓	✓	-	-
Dell Technologies	✓	✓	✓	✓	✓	✓	-	✓	-	✓
DigitalOcean	✓	-	-	✓	-	-	-	-	✓	-
Everpure	-	✓	✓	✓	✓	-	✓	✓	✓	✓

	TARGET MARKET				DEPLOYMENT MODEL					
PLATFORM	SMB	LARGE ENTERPRISE	CSP	MSP	PHYSICAL APPLIANCE	VIRTUAL APPLIANCE	PUBLIC CLOUD IMAGE	SOFTWARE ONLY	SAAS	SELF MANAGED
Hitachi Vantara	✓	✓	✓	✓	✓	✓	-	✓	✓	✓
IBM	-	✓	✓	✓	✓	✓	-	✓	✓	✓
Infinidat	-	✓	✓	✓	✓	-	✓	-	-	✓
MinIO	✓	✓	✓	✓	✓	-	✓	✓	-	✓
NetApp	-	✓	✓	✓	✓	✓	-	✓	-	✓
Nutanix	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
OSNexus	✓	✓	✓	✓	✓	✓	✓	✓	-	✓
Quantum	-	✓	✓	✓	✓	✓	-	✓	✓	✓
Qumulo	-	✓	✓	✓	✓	✓	✓	✓	✓	✓
Quobyte	-	✓	✓	✓	✓	✓	✓	✓	-	✓
Scality	✓	✓	✓	✓	-	-	-	✓	✓	✓
Seagate	✓	✓	✓	✓	-	-	-	-	✓	-
SoftIron	-	✓	✓	✓	✓	-	-	-	-	✓
Spectra Logic	-	✓	✓	✓	✓	✓	✓	✓	-	✓
Storj	✓	✓	✓	✓	-	-	-	-	✓	✓
StorONE	✓	✓	✓	✓	✓	✓	✓	✓	-	-
VAST Data	-	✓	✓	✓	✓	-	✓	✓	✓	✓
WEKA	-	✓	✓	✓	✓	✓	✓	✓	✓	-
Zadara	✓	✓	✓	✓	-	✓	✓	-	✓	-

Source: GigaOm 2026

Table 1 components are evaluated in a binary yes/no manner and do not factor into a vendor's designation as a Leader, Challenger, or Entrant on the Radar chart (**Figure 1**).

"Target market" reflects which use cases each solution is recommended for, not simply whether that group can use it. For example, if an SMB could use a solution but doing so would be cost-prohibitive, that solution would be rated "no" for SMBs.

03 Decision Criteria Comparison

ALL SOLUTIONS INCLUDED in this Radar report meet the following table stakes—capabilities widely adopted and well implemented in the sector:

- Access controls
- Indexing and search
- Data protection and optimization
- Encryption and security
- Remote and geo-replication
- Scale-out architecture
- S3 and other protocols

Tables 2, 3, and 4 summarize how each vendor in this research performs in the areas we consider differentiating and critical in this sector. The objective is to give the reader a snapshot of the technical capabilities of available solutions, define the perimeter of the relevant market space, and gauge the potential impact on the business.

- Key features differentiate solutions, highlighting the primary criteria to be considered when evaluating an object storage solution
- Emerging features show how well each vendor implements capabilities that are not yet mainstream but are expected to become more widespread and compelling within the next 12 to 18 months
- Business criteria provide insight into the nonfunctional requirements that factor into a purchase decision and determine a solution's impact on an organization

These decision criteria are summarized on the next page.

Key Features:

- **Kubernetes support:** This evaluates the depth of integration with container orchestration platforms. Leaders provide native Container Storage Interface (CSI) drivers and Operators for dynamic provisioning, while advanced support includes the emerging Container Object Storage Interface (COSI) standard to manage buckets as first-class Kubernetes resources.
- **Workload optimization:** This criterion assesses how well the solution tunes performance for specific use cases, such as AI/ML training or high-frequency analytics. Key capabilities include QoS controls, parallel data loading (like S3 over RDMA), and flash-optimized I/O paths that maximize throughput.
- **Auditing:** Robust auditing requires granular logging of all management and data access events. Top solutions offer native integration with SIEM platforms (like Splunk) and provide tamper-proof audit trails to satisfy strict regulatory compliance and forensic analysis requirements.
- **Versioning:** This feature ensures data integrity by maintaining multiple variants of an object in the same bucket. It is the foundational mechanism for recovery from accidental deletion or overwrite and is a prerequisite for implementing effective lifecycle management and immutability policies.
- **Ransomware protection:** Beyond basic backups, this criterion focuses on immutable storage capabilities, specifically S3 Object Lock (Compliance Mode). Advanced solutions also incorporate active defenses, such as AI-driven anomaly detection, to identify and block malicious encryption activity or mass deletions in real time.
- **Reporting and analytics:** Solutions must provide deep visibility into system health, capacity usage, and performance trends. Leaders go beyond basic dashboards to offer predictive analytics for capacity planning and granular insights into data consumption patterns, often via integrated tools or open metrics standards like Prometheus.
- **Storage optimization:** This measures the efficiency of data reduction technologies, including erasure coding, compression, and deduplication. High scores are awarded to architectures that maximize usable capacity and minimize hardware footprint without compromising read/write performance or data durability.
- **Public cloud integration:** This capability assesses how well the solution supports hybrid cloud workflows. It includes features like bidirectional replication, automated tiering to public cloud targets (for example, AWS S3, Azure Blob), and the ability to present a unified global namespace across on-prem and cloud environments.

- **Global namespace:** This architecture federates distinct storage clusters across multiple locations into a single, logical view. It creates a unified endpoint that abstracts physical data placement, allowing applications to read and write data anywhere while the system manages active-active consistency and geo-replication in the background.

Table 2. Key Features Comparison

KEY FEATURES	AVERAGE SCORE	KUBERNETES SUPPORT	WORKLOAD OPTIMIZATION	AUDITING	VERSIONING	RANSOMWARE PROTECTION	REPORTING & ANALYTICS	STORAGE OPTIMIZATION	PUBLIC CLOUD INTEGRATION	GLOBAL NAMESPACE
Cloudian	4.0	***	*****	****	*****	****	***	***	*****	****
Cohesity	3.4	**	***	****	***	*****	****	****	***	***
Cubbit	3.3	*	***	**	****	****	***	****	****	*****
DataCore Software	3.4	***	****	****	****	***	***	***	***	****
DDN	4.2	****	*****	****	****	****	****	*****	****	****
Dell Technologies	4.3	*****	****	*****	****	****	*****	****	****	****
DigitalOcean	2.2	**	***	***	***	*	**	**	***	*
Everpure	4.0	*****	****	***	****	*****	*****	***	****	***
Hitachi Vantara	4.0	***	****	****	****	****	****	*****	****	****
IBM	4.2	*****	****	***	****	*****	****	*****	****	****
Infinidat	4.2	****	****	*****	*****	****	****	*****	****	***
MinIO	4.0	*****	*****	****	*****	***	***	***	****	****
NetApp	4.4	*****	****	*****	*****	****	****	***	*****	*****
Nutanix	4.6	****	*****	*****	*****	****	*****	****	*****	****
OSNexus	3.7	***	****	***	****	****	****	***	****	****
Quantum	3.4	***	****	***	****	****	***	***	****	***

KEY FEATURES	AVERAGE SCORE	KUBERNETES SUPPORT	WORKLOAD OPTIMIZATION	AUDITING	VERSIONING	RANSOMWARE PROTECTION	REPORTING & ANALYTICS	STORAGE OPTIMIZATION	PUBLIC CLOUD INTEGRATION	GLOBAL NAMESPACE
Qumulo	3.9	****	***	****	***	****	*****	***	*****	****
Quobyte	3.8	****	****	***	****	****	****	****	***	****
Scality	4.1	****	****	****	***	*****	****	****	****	*****
Seagate	2.9	**	***	****	***	***	***	**	*****	*
SoftIron	3.0	***	***	****	***	****	***	***	**	**
Spectra Logic	3.0	**	***	***	***	****	***	***	***	***
Storj	2.9	****	**	***	****	***	**	**	***	***
StorONE	3.0	**	****	***	***	***	***	***	***	***
VAST Data	4.1	***	*****	****	***	****	****	*****	****	*****
WEKA	3.1	****	****	***	–	***	***	*****	*****	*
Zadara	3.4	***	***	****	***	****	***	***	*****	***

***** Exceptional **** Superior *** Capable ** Limited * Poor – Not applicable or absent

Source: GigaOm 2026

Emerging Features

- **Object content indexing:** This capability extends beyond basic metadata search to allow deep inspection and querying of object payloads (JSON, CSV, Parquet, and the like), using SQL-like syntax (S3 Select). It enables applications to filter data server-side, significantly reducing network traffic and accelerating analytics workflows.
- **COSI support:** The COSI is an emerging Kubernetes standard designed to provision and manage object storage buckets as native resources. Unlike standard CSI drivers, COSI provides a vendor-neutral API for bucket lifecycle management, ensuring true portability for cloud-native applications.
- **AI-optimized data path:** To meet the demands of AI/ML training, this feature optimizes the I/O path to bypass legacy kernel bottlenecks. Key technologies include S3 over RDMA and NVIDIA GPUDirect Storage support, which facilitate direct data transfers from storage to GPU memory, maximizing throughput and reducing latency.

Table 3. Emerging Features Comparison

EMERGING FEATURES	AVERAGE SCORE	OBJECT CONTENT INDEXING	COSI SUPPORT	AI-OPTIMIZED DATA PATH
Cloudian	3.7	****	***	****
Cohesity	3.0	*****	*	***
Cubbit	0.7	*	-	*
DataCore Software	2.0	***	-	***
DDN	2.0	***	-	***
Dell Technologies	2.7	***	*	****
DigitalOcean	0.7	*	-	*
Everpure	3.3	*	****	*****
Hitachi Vantara	2.3	***	-	****
IBM	3.3	*****	**	***
Infinidat	2.3	***	*	***
MinIO	4.0	*****	***	****
NetApp	2.3	***	*	***
Nutanix	4.0	***	*****	****
OSNexus	0.7	*	-	*
Quantum	1.3	*	-	***
Qumulo	1.0	*	-	**
Quobyte	1.7	*	-	****
Scality	4.0	***	*****	****
Seagate	2.3	***	*	***
SoftIron	1.0	-	-	***
Spectra Logic	1.3	***	-	*
Storj	1.3	-	*	***

EMERGING FEATURES	AVERAGE SCORE	OBJECT CONTENT INDEXING	COSI SUPPORT	AI-OPTIMIZED DATA PATH
StorONE	2.0	***	-	***
VAST Data	4.3	****	****	*****
WEKA	2.7	***	-	*****
Zadara	1.0	*	-	***

***** Exceptional **** Superior *** Capable ** Limited * Poor – Not applicable or absent

Source: GigaOm 2026

Business Criteria

- **Cost transparency:** This criterion evaluates the predictability of the pricing model. Leaders avoid bill shock by offering simple, capacity-based licensing or flat-rate cloud pricing that eliminates complex metering for API requests and data egress, ensuring TCO remains clear and aligned with budget expectations.
- **Performance:** Once relegated to "cheap and deep" archives, modern object storage must deliver high throughput and low latency. This score reflects the solution's ability to saturate network links (100GbE+) and support demanding transactional workloads, such as AI training or high-speed analytics, often leveraging all-flash architectures.
- **Flexibility:** This measures the versatility of the solution across deployment models (appliance, software-defined, SaaS) and protocols. High scores go to platforms that allow organizations to avoid vendor lock-in, mix hardware generations, and unify access via multiprotocol support—S3, Network File System (NFS), Server Message Block (SMB)—on a single dataset.
- **Manageability:** Efficient operations are critical at petabyte scale. This criterion assesses the quality of the management plane, prioritizing zero-touch automation, AIOps-driven health monitoring, and centralized control consoles that simplify day-to-day tasks like provisioning, upgrades, and troubleshooting across distributed clusters.
- **Scalability:** The defining characteristic of object storage is its ability to grow. This evaluates the architecture's capacity to scale nondisruptively from terabytes to exabytes and billions of objects. Leading vendors demonstrate linear performance scaling and the ability to expand without introducing sharding complexity or performance bottlenecks.
- **Ecosystem:** A robust ecosystem extends beyond validated independent software vendor (ISV) integrations (backup, AI, analytics) to include the strength of the user community and knowledge base. High scores reflect active user groups, accessible training, and certification programs that ensure staff can effectively support and scale the platform.

Table 4. Business Criteria Comparison

BUSINESS CRITERIA COMPARISON	AVERAGE SCORE	COST TRANSPARENCY	PERFORMANCE	FLEXIBILITY	MANAGEABILITY	SCALABILITY	ECOSYSTEM
Cloudian	3.2	***	****	***	***	***	***
Cohesity	3.7	****	**	***	****	*****	****
Cubbit	3.5	****	***	***	****	****	***
DataCore Software	3.0	***	***	***	***	***	***
DDN	4.0	***	*****	***	****	*****	****
Dell Technologies	3.7	***	****	***	***	****	*****
DigitalOcean	3.5	*****	***	***	****	***	***
Everpure	4.3	***	*****	*****	*****	****	****
Hitachi Vantara	3.8	****	****	****	***	****	****
IBM	3.8	***	****	***	****	*****	****
Infinidat	3.5	****	***	***	***	****	****
MinIO	4.2	****	*****	***	***	*****	*****
NetApp	3.8	***	***	***	*****	****	*****
Nutanix	3.5	***	***	****	****	****	***
OSNexus	3.3	***	***	****	***	****	***
Quantum	3.8	****	***	****	*****	****	***
Qumulo	3.5	***	***	****	****	****	***
Quobyte	3.5	***	***	****	****	****	***
Scality	3.8	***	****	****	****	*****	***
Seagate	3.5	*****	***	***	***	****	***

BUSINESS CRITERIA COMPARISON	AVERAGE SCORE	COST TRANSPARENCY	PERFORMANCE	FLEXIBILITY	MANAGEABILITY	SCALABILITY	ECOSYSTEM
SoftIron	3.5	***	***	*****	****	***	***
Spectra Logic	3.2	***	****	***	***	***	***
Storj	3.3	***	****	***	***	****	***
StorONE	3.3	***	****	****	***	***	***
VAST Data	4.5	****	*****	*****	****	*****	****
WEKA	4.3	***	*****	*****	****	*****	****
Zadara	4.2	*****	****	****	*****	****	***

***** Exceptional **** Superior *** Capable ** Limited * Poor – Not applicable or absent

Source: GigaOm 2026

Business criteria provide insight into the nonfunctional requirements that factor into a purchase decision and determine a solution's impact on an organization.

04 GigaOm Radar

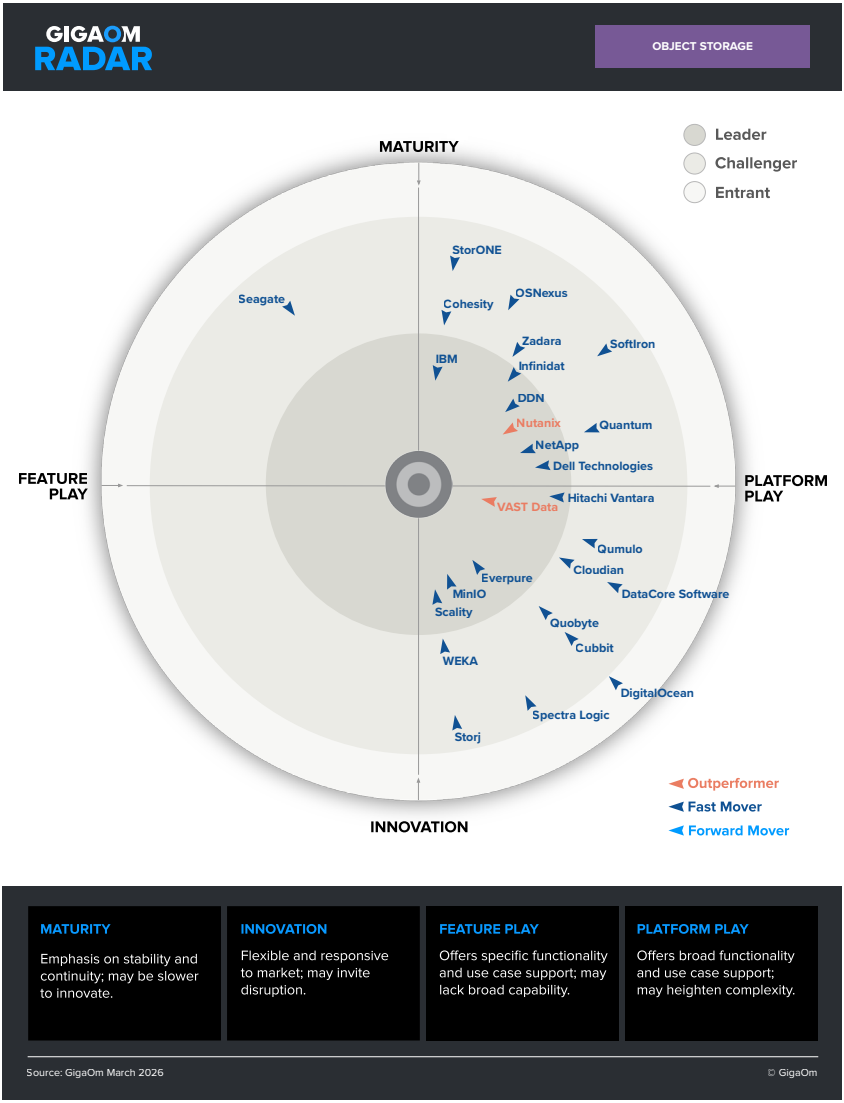

Figure 1. GigaOm Radar for Object Storage v7

THE GIGAOM RADAR PLOTS VENDOR SOLUTIONS across a series of concentric rings with those set closer to the center judged to have the most complete solutions. The chart characterizes each vendor on two axes—balancing Maturity versus Innovation and Feature Play versus Platform Play—while providing an arrowhead that projects each solution's expected evolution over the coming 12 to 18 months.

As you can see in **Figure 1**, this year's Radar for object storage reflects a market in active transformation. The field has expanded from 22 to 27 evaluated solutions, with five new vendors joining the report, yet only a single Entrant appears on the chart, which is a clear signal that this is a competitive space where meaningful participation demands enterprise-grade capabilities from day one.

The most striking shift from last year is the impact of agentic AI and AIOps on vendor roadmaps. Whereas the v6 Radar showed a market gravitating toward stability and incremental refinement, v7 reveals a wave of disruption as vendors race to embed AI-driven automation, intelligent tiering, and autonomous operations into their object storage platforms. This urgency has pulled several solutions toward the Innovation hemisphere. Their products look and feel materially different year over year as they respond to enterprise demand for AI-optimized data infrastructure. The result is a more balanced distribution across the Maturity/Innovation axis than we have seen in prior iterations, even as the Maturity hemisphere retains a strong cluster of vendors that continue to prioritize consistency, compatibility, and incremental improvement.

The Platform Play hemisphere dominates decisively. Nearly every evaluated solution has expanded beyond core object storage to offer unified data platforms, multiprotocol access, integrated analytics pipelines, and AI/ML workflow support. A single vendor stands as the lone Feature Play, a deliberate positioning that reflects its focused, storage-as-a-service approach rather than the broad platform ambitions of its peers. This near-universal drift toward Platform Plays underscores a market reality: enterprises no longer evaluate object storage in isolation. They expect a single platform to serve as the foundation for backup, archive, analytics, and, increasingly, AI training and inference workloads.

The Leaders ring is densely populated, with roughly a dozen solutions achieving top-tier positioning. This concentration reflects the overall maturity of the object storage market, in which many vendors now deliver robust, enterprise-grade capabilities across the evaluation criteria. However, the number of Outperformers has contracted sharply, dropping from five in v6 to just two. This compression tells an important story: the competitive field has tightened, and the pace of differentiation that propelled several vendors to Outperformer status last year has leveled as AI-driven capabilities become table stakes rather than differentiators. The two remaining Outperformers have sustained exceptional forward momentum, but the broader field has narrowed the gap.

Year-over-year vendor movement reinforces these dynamics. Several solutions that held Outperformer status in v6 have settled into Fast Mover trajectories, still advancing but at a pace more consistent with the broader market. Meanwhile, a few vendors have made notable upward moves into the Leaders ring, demonstrating that strong execution on AI integration, platform unification, and operational simplicity can still break through in a

crowded field. Conversely, at least one former Leader has pulled back to Challenger status, a reminder that standing still in a market this dynamic carries real consequences.

The trajectory arrows across the chart point predominantly inward and toward Innovation, suggesting that over the next 12 to 18 months, vendors will continue to invest aggressively in AI-powered operations, autonomous data management, and performance optimization for GPU-driven workloads. Organizations evaluating object storage should pay close attention to whether a vendor's innovation trajectory aligns with their own data strategy, particularly around AI readiness and operational automation.

In reviewing solutions, it's important to keep in mind that there are no universal "best" or "worst" offerings; every solution has aspects that might make it a better or worse fit for specific customer requirements. Prospective customers should consider their current and future needs when comparing solutions and vendor roadmaps.

INSIDE THE GIGAOM RADAR

To create the GigaOm Radar graphic, key features, emerging features, and business criteria are scored and weighted. Key features and business criteria receive the highest weighting and have the most impact on vendor positioning on the Radar graphic. Emerging features receive a lower weighting and have a lower impact on vendor positioning on the Radar graphic. The resulting chart is a forward-looking perspective on all the vendors in this report, based on their products' technical capabilities and roadmaps.

Note that the Radar is technology-focused, and business considerations such as vendor market share, customer share, spend, recency or longevity in the market, and so on are not considered in our evaluations. As such, these factors do not impact scoring and positioning on the Radar graphic.

For more information, please visit our Methodology.

05 Solution Insights

Cloudian: HyperStore

SOLUTION OVERVIEW

Cloudian specializes in enterprise-grade object storage software and appliances, focusing on S3 API compatibility and modular scalability. The company targets large enterprises and service providers, offering a software-defined architecture that runs on standard x86 hardware or proprietary appliances. Recent strategic shifts have moved the platform beyond secondary storage into primary workloads, as evidenced by the introduction of the HyperScale AI Data Platform and support for NVIDIA GPUDirect Storage.

HyperStore is designed not as a niche repository but as a comprehensive foundation for diverse data workflows, including data lakes, backup targets, and AI training sets. The solution will look and feel different over the contract lifecycle. Cloudian delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. Primary emphasis is placed on AI infrastructure integration and performance optimization to address high-throughput demands.

Cloudian is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Cloudian scored well on a number of the decision criteria, including:

- **Workload optimization:** The solution aggressively targets AI and ML workloads through its HyperScale AI Data Platform, which integrates compute and storage to minimize latency. Administrators can use granular QoS controls to manage bandwidth per user or group, ensuring critical tasks receive necessary resources. The system also employs intelligent tiering policies that automatically move data between flash and HDD tiers based on access patterns.
- **Versioning:** HyperStore supports full S3 object versioning, enabling users to preserve, retrieve, and restore every version of an object stored. This capability integrates tightly with lifecycle management policies, allowing for the automated expiration or tiering of older versions to control capacity costs. It serves as the foundation for the system's immutability features, ensuring that locked versions remain protected against modification or deletion.

- **Ransomware protection:** Cloudian provides a hardened security profile with its S3 Object Lock implementation, supporting both compliance and governance modes to prevent data deletion. The solution includes rigorous system hardening measures, such as a "Root Disable" feature that prevents even administrators with physical access from compromising immutability settings. Secure boot processes further defend against OS-level compromises, satisfying strict federal requirements.

OPPORTUNITIES

Cloudian has room for improvement in a few decision criteria, including:

- **Auditing:** HyperStore generates detailed logs of system and data access events, which are essential for compliance and security monitoring. However, while the system supports comprehensive logging, direct real-time streaming to third-party SIEM tools typically relies on standard syslog forwarding rather than deep, native API integration found in some security-specialized tools.
- **Reporting and analytics:** The HyperIQ platform provides comprehensive monitoring, predictive analytics for capacity planning, and user behavior analytics to detect anomalies. While these capabilities are extensive, they require the deployment and management of the separate HyperIQ tool to fully visualize access patterns and track user behavior rather than these advanced insights being embedded directly in the core storage management console.
- **Storage optimization:** The platform uses efficient erasure coding schemes and supports multiple compression algorithms (lz4, snappy, zlib) to minimize storage overhead. However, HyperStore lacks a global always-on inline deduplication engine for all data types; instead, it typically relies on integration with backup partners like Veeam or Rubrik to achieve deduplication efficiency.

PURCHASE CONSIDERATIONS

Cloudian uses a transparent capacity-based subscription model that avoids the hidden costs often associated with public cloud storage. Specifically, the vendor does not charge for data ingress, egress, or API calls, which provides significant budget predictability for organizations with active data workflows. This licensing structure is particularly advantageous for repatriation projects whose unpredictable public cloud fees are a driver.

The solution is highly flexible in deployment, offered as software-only for bare metal, as preconfigured appliances, or as virtual machines for nonproduction testing. This software-defined approach allows buyers to decouple software licensing from hardware procurement,

leveraging existing server vendor relationships to lower TCO. The architecture is designed for scale, making it suitable for large enterprises and MSPs requiring multitenancy, rather than small-scale deployments where the overhead of a distributed ring architecture might be excessive.

USE CASES

Cloudian supports highly regulated verticals, including federal government, healthcare, and finance, by maintaining certifications such as FIPS 140-2/3 and SEC Rule 17a-4(f). The platform is well suited for AI/ML infrastructure, utilizing RDMA and GPUDirect Storage to feed data directly to GPUs for high-performance training. Additionally, it serves as a robust target for ransomware protection workflows, offering an immutable air gap for backup data.

Cohesity: SmartFiles*

SOLUTION OVERVIEW

Cohesity targets the midsize to large enterprise and service provider sectors with a focus on consolidating disparate data silos: backup, file, object, dev/test, and analytics onto a single web-scale fabric. Originally established in the secondary storage and backup arena, the vendor has evolved its platform into a comprehensive Data Cloud designed for data management, security, and insight rather than just raw storage capacity. This strategic shift emphasizes the value of data beyond simple retention, integrating security and AI-driven search capabilities directly into the storage layer to address mass data fragmentation.

The core solution, Cohesity SmartFiles, is a software-defined, scale-out file and object service that runs on the Cohesity Data Cloud platform. It is available as physical appliances through hardware partners, as virtual appliances for edge deployments, or as a cloud-native edition on major hyperscalers. SmartFiles offers a unified multiprotocol experience, allowing simultaneous access to data via NFS, SMB, and S3 APIs. The solution creates a converged environment where data protection and file/object services coexist, reducing the need for separate infrastructure silos.

The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. Cohesity values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement.

Cohesity is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Cohesity scored well on a number of decision criteria, including:

- **Storage optimization:** Cohesity employs global variable-length deduplication and compression across the entire cluster and all workloads, including files, objects, and backups. This approach identifies duplicate data blocks even when shifted within a file, delivering significant data reduction ratios. The architecture supports erasure coding and small file optimization to further reduce the physical storage footprint, optimizing TCO for capacity-heavy environments.
- **Auditing:** The platform provides granular, real-time logging of file and object access events, including permission changes, which are actionable rather than passive. Through the Helios management plane, these logs fuel AI-driven anomaly detection to identify potential security threats like ransomware. Additionally, integrations with SIEM platforms allow for the streaming of audit data to external security tools for broader correlation and forensic analysis.
- **Ransomware protection:** SmartFiles incorporates a defense-in-depth strategy centered on immutable storage via DataLock, which enforces time-bound write once, read many (WORM) protection to prevent unauthorized modification or deletion. This is complemented by multifactor authentication and quorum controls for sensitive administrative actions. The DataHawk service further enhances resilience by actively monitoring data ingest patterns for anomalies, such as high-entropy writes or massive deletions, to trigger automated alerts.

OPPORTUNITIES

Cohesity has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While the platform offers robust protection for Kubernetes via CSI snapshots for backup and recovery, it currently lacks a native CSI driver for dynamic provisioning of primary object storage. Applications in Kubernetes environments typically access Cohesity storage through manual S3 API connections or external drivers rather than via seamless, native kubectl provisioning commands, limiting its utility as a primary storage backend for containerized workloads.
- **Workload optimization:** The solution features automated tiering to move cold data to public cloud targets and basic QoS controls to manage performance prioritization. However, the general-purpose architectural focus means it lacks the advanced, parallel I/O layout optimizations found in specialized high-performance computing (HPC) storage systems. It is designed primarily for consolidation and management rather than the extreme low-latency requirements of dedicated AI training scratchpads.

- **Versioning:** SmartFiles supports standard S3 bucket versioning and lifecycle management policies, which allow users to recover from accidental overwrites and manage object expiration for compliance. However, the implementation is focused on API compatibility and basic retention, lacking the hypergranular controls found in some cloud-native competitors, such as complex branching of object histories or per-object MFA-Delete policies outside of global settings.

PURCHASE CONSIDERATIONS

Cohesity's licensing model is transparent and increasingly flexible, addressing the common industry challenge of unpredictable costs. The introduction of a pay-per-use consumption model allows organizations, particularly MSPs, to align infrastructure spending directly with usage rather than committing to massive upfront capital expenditures. This model, combined with high storage efficiency, provides a predictable cost structure. The software is available on certified hardware from major vendors like HPE, Cisco, and Dell, or as a software-only deployment, giving buyers a significant choice in their supply chain.

The solution is best suited for large enterprises and service providers grappling with petabyte-scale data growth and fragmentation. Deployment is generally straightforward due to the appliance-based model and the unified Helios management interface, which simplifies the configuration of global clusters. However, because SmartFiles is often part of a broader platform consolidation strategy involving backups and archives, migration planning can be complex compared to deploying a simple, standalone storage array. Organizations looking for a quick, tactical fix for a single workload may find the platform's comprehensive nature more functionality than required, whereas those seeking to replace legacy NAS and backup targets simultaneously will find the most value.

USE CASES

Cohesity supports a broad range of industry verticals, including financial services, healthcare, and government because data compliance and security are paramount for them. Specific use cases focus on the consolidation of secondary data, including active archives, file shares, and backup repositories, into a single searchable namespace. The platform is particularly effective for organizations needing to eliminate dark data by leveraging integrated analytics and search to derive value from dormant assets. It also serves hybrid cloud use cases well, enabling seamless data mobility and disaster recovery across on-prem clusters and public cloud infrastructure.

Cubbit: DS3 Composer

SOLUTION OVERVIEW

Cubbit is a European software vendor focused on distributed, sovereign data management architectures. The company distinguishes itself through a peer-to-peer design that aggregates fragmented storage resources into encrypted, geo-distributed networks. The primary offering, DS3 Composer, is a software-defined object storage platform consisting of three main components: a centralized Coordinator (available as SaaS or self-hosted) for metadata and orchestration, storage Agents that run on customer-provided hardware (bare metal, VMs, or public cloud instances), and S3-compatible Gateways. This architecture allows MSPs and enterprises to build private, S3-compatible cloud storage regions using their own infrastructure or federated resources.

The solution will look and feel different over the contract lifecycle. Cubbit delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. Primary emphasis is on emerging features and rapid development to address gaps.

Cubbit is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Cubbit scored well on a number of decision criteria, including:

- **Versioning:** The platform provides a comprehensive implementation of S3 bucket versioning that adheres strictly to AWS standards. It supports both compliance and governance modes, enabling the recovery of overwritten or deleted objects. This capability serves as a reliable foundation for lifecycle management rules and integrates seamlessly with Object Lock to ensure data immutability for regulatory compliance.
- **Storage optimization:** Cubbit employs advanced Reed-Solomon erasure coding to achieve high data durability with significantly lower storage overhead compared to traditional replication. By splitting encrypted data into shards and distributing them across the network, the system achieves durability levels of up to 15 nines while maintaining a redundancy ratio between 1.5x and 1.8x. This approach decouples resilience from hardware reliability, allowing for the efficient use of heterogeneous or older storage media.
- **Ransomware protection:** Defense against data compromise is multilayered, combining S3 Object Lock with an optional customer-enabled zero-knowledge encryption architecture. Data is encrypted client-side using AES-256 before ingestion, ensuring neither the storage provider nor the infrastructure operators can access

the raw data. This immutability, validated by backup ecosystem certifications, provides a robust passive defense against encryption attacks and unauthorized deletion.

OPPORTUNITIES

Cubbit has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While the vendor's own infrastructure components are deployed as containers within Kubernetes clusters, the solution currently lacks a production-ready CSI driver. This limits the ability of developers to dynamically provision persistent volumes for stateful applications directly through standard Kubernetes manifests, restricting the platform's utility in modern cloud-native DevOps pipelines.
- **Workload optimization:** The architecture is well suited for capacity-driven workloads like backup and active archive but faces limitations in high-performance transactional scenarios. The network latency inherent in wide-area geo-distribution and shard reconstruction makes the solution less suitable for latency-sensitive databases or HPC compared to local flash arrays. Additionally, specialized data paths for AI, such as GPUDirect, are not yet supported.
- **Auditing:** The platform currently relies on basic logging to files or local databases for tracking system operations rather than offering native push-based integration with major SIEM systems. While performance metrics are well exposed via Prometheus, the lack of automated audit log forwarding to enterprise security tools increases operational friction for organizations with strict compliance reporting requirements.

PURCHASE CONSIDERATIONS

Cubbit differentiates itself with a radically transparent licensing model designed to eliminate the unpredictability of cloud costs. The vendor charges a single flat rate per terabyte per month, with no hidden fees for egress, API requests, or bucket deletion. This structure is particularly attractive to MSPs and organizations managing data-intensive restore operations that would incur significant penalties with hyperscalers.

The solution is primarily targeted at MSPs looking to white label storage services and European enterprises with strict data sovereignty mandates. Deployment is flexible, supporting a software-only model that runs on commodity hardware, including repurposed trashware or a hybrid model using the SaaS control plane. While the SaaS Coordinator simplifies day-to-day management by abstracting the complexity of the distributed network, enabling the optional

zero-knowledge architecture imposes a requirement for rigorous client-side key management, as the vendor cannot recover lost credentials. Migration is facilitated by standard S3 API compatibility, though legacy file protocol support is gateway-based rather than native.

USE CASES

Cubbit focuses heavily on the MSP and critical infrastructure sectors, including government and defense, for which digital sovereignty is paramount. The solution is ideally suited for cold and warm data use cases such as immutable backup targets, validated by integrations with Commvault, Veeam, and Acronis, and active archiving. The geo-distributed nature of the swarm architecture also supports edge content delivery scenarios, allowing data to be cached and accessed closer to users without the complexity of managing distinct replication links.

DataCore Software: Swarm

SOLUTION OVERVIEW

DataCore Software specializes in SDS, abstracting control from the physical layer to aggregate disparate hardware resources into unified, resilient pools. Swarm is the vendor's flagship object storage platform, engineered as a massively scalable, hardware-agnostic solution that decouples storage intelligence from underlying media. It operates as a core pillar of the broader DataCore.NEXT vision, which aims to unify block, file, and object storage into a cohesive fabric. Swarm employs a parallelized architecture designed to function either as a primary storage tier for high-performance workloads or as a deep archive while eliminating the constraints of proprietary hardware lock-in.

The solution will look and feel different over the contract lifecycle. DataCore Software delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. Primary emphasis is placed on emerging edge capabilities and container-native integrations.

DataCore Software is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

DataCore Software scored well on a number of decision criteria, including:

- **Versioning:** Swarm implements standard S3 object versioning, allowing users to list, retrieve, and restore previous iterations of an object. This capability integrates closely with data protection

strategies, preserving writes to create a history of changes for point-in-time recovery. The architecture ensures that even if a delete marker is placed on an object, the underlying versions remain preserved and protected, facilitating reliable recovery workflows.

- **Storage optimization:** The platform achieves high capacity efficiency, claiming up to 95% utilization of available drive capacity through advanced erasure coding implementation. Swarm boots directly from RAM, leaving nearly the entire disk volume available for data payload. Additionally, the Darkive power conservation technology optimizes energy costs by spinning down or powering off disks during periods of inactivity, addressing sustainability and cost concerns at scale.
- **Ransomware protection:** Swarm supports S3 Object Lock in both governance and compliance modes, creating a virtual air gap against credential compromise. In compliance mode, data remains immutable until retention periods expire. This feature is reinforced by DataCore Software's proprietary Integrity Seals, which use cryptographic hashing to verify bit-level authenticity, ensuring data has not been tampered with.

OPPORTUNITIES

DataCore Software has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While DataCore Software supports Kubernetes via its Puls8 container-native storage solution and CSI driver, Swarm lacks deep native integration for object management within the Kubernetes control plane. The solution currently relies on separate provisioning layers rather than employing a specialized object operator or custom resource definitions (CRDs) to manage buckets and users natively.
- **Reporting and analytics:** The platform includes embedded tools for metering usage by each tenant and monitoring cluster health, but it does not offer intrinsic AI-driven predictive analytics within the base Swarm codebase. Advanced forecasting for capacity planning or drive failure prediction typically requires integration with external tools like Grafana rather than being a fully autonomous on-box capability.
- **Public cloud integration:** Swarm facilitates tiering and replication to S3-compatible public cloud targets, effectively extending on-prem capacity. However, this integration focuses primarily on unidirectional data movement to the cloud. It does not currently offer a fully unified, bidirectional global namespace that treats on-prem and cloud buckets as a single logical entity with seamless write-anywhere capabilities.

PURCHASE CONSIDERATIONS

DataCore Software differentiates itself with a transparent, capacity-based licensing model that charges for usable storage rather than raw capacity or node count. This approach eliminates the hidden costs often associated with public cloud storage, such as API transaction and egress fees, making TCO highly predictable. The software-defined nature of Swarm allows organizations to deploy on standard x86 servers from any vendor, mixing hardware generations within the same cluster to avoid vendor lock-in and optimize supply chain costs. It supports diverse deployment models, including physical appliances via partners, virtual machines, and software-only installations on standard x86 hardware that boot from RAM. This flexibility makes enterprise-grade object storage accessible to a wide range of buyers, from SMBs starting with single-node clusters to large enterprises managing exabytes.

USE CASES

Swarm serves a broad spectrum of industry verticals, including healthcare for vendor neutral archives (VNA) and Picture Archiving and Communication System (PACS) imaging, and media and entertainment for high-throughput video on demand (VOD) streaming and origin storage. Its architecture is particularly well suited for active archive scenarios where massive datasets such as legal records or digital assets must be retained cost-effectively but remain instantly accessible via HTTP/S3. Additionally, Swarm acts as a secure, immutable target for backup and ransomware protection workflows, validated by certifications with major vendors like Veeam and Commvault.

DDN: Infinia

SOLUTION OVERVIEW

DataDirect Networks (DDN) focuses on high-performance storage solutions for AI, HPC, and multicloud environments. Leveraging decades of experience powering supercomputers, DDN addresses the distinct data management challenges of large-scale deep learning and generative AI workflows.

Infinia is a software-defined, hyper-converged object storage platform designed as a distributed key-value (KV) store rather than a traditional file-system overlay. Part of DDN's broader Data Intelligence Platform, Infinia is engineered specifically for consolidating massive unstructured datasets across edge, core, and cloud locations. It is available as a standalone software solution, a turnkey appliance (AI400X2 or AI2200), or via public cloud marketplaces.

DDN employs a platform strategy focused on stability and operational continuity for enterprise AI infrastructure. The solution prioritizes consistent user experiences and assured compatibility to support mission-critical model training and inference pipelines over experimental feature releases.

DDN is positioned as a Leader and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

DDN scored well on a number of decision criteria, including:

- **Workload optimization:** Infinia is explicitly tuned for AI and ML pipelines, addressing specific bottlenecks like GPU starvation. It utilizes a B-EpsilonTree data structure to efficiently pack small objects, which are common in training datasets, eliminating the block alignment waste typical of traditional file systems. This ground-up architecture is designed as a distributed key-value store rather than a protocol overlay, optimizing key-value cache retrieval to reduce time to first token (TTFT) for large language model (LLM) inference, ensuring expensive GPU resources remain saturated.
- **Storage optimization:** The platform maximizes effective capacity through inline, object-level compression that dynamically selects algorithms (Zstd, LZ4) based on data characteristics. The Adaptive I/O Engine intelligently routes data, using replication for small latency-sensitive writes and erasure coding for large sequential streams to balance performance with storage efficiency. This approach supports high-density storage suitable for massive data lakes without performance degradation.
- **Ransomware protection:** DDN employs a comprehensive defense-in-depth strategy centered on immutability and anomaly detection. Infinia supports S3 Object Lock in both governance and compliance modes, preventing the deletion or modification of critical data for regulatory adherence. This is augmented by telemetry-based monitoring that analyzes more than 1,000 metrics per node to detect behavioral anomalies, such as sudden spikes in encryption or deletion activity, that often signal a ransomware attack.

OPPORTUNITIES

DDN has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While Infinia provides a robust CSI driver for dynamic volume provisioning, it lacks support for the emerging COSI standard. Furthermore, while ReadWriteOnce (RWO) support is mature, ReadWriteMany (RWX) capabilities for block storage are

currently limited, which may restrict certain shared-volume use cases in complex microservices environments.

- **Reporting and analytics:** The platform captures detailed telemetry but relies heavily on external tools for visualization and advanced analysis. While logs can be exported via OpenTelemetry to systems like Splunk or Datadog, the native interface lacks deep, built-in visual analytics and content indexing capabilities compared to some competitors. Additionally, log retention policies are currently fixed based on severity levels rather than being fully user-configurable for long-term audit trails within the platform itself.
- **Public cloud integration:** Infinia is available on Google Cloud Platform and Oracle Cloud Infrastructure, but its integration with specific hyperscaler ecosystems has limitations. It does not currently support niche AWS-specific features such as S3 Glacier Instant Retrieval transitions or intelligent tiering tied to AWS native storage classes. Furthermore, achieving true active-active consistency for the same object across widely separated geographic regions remains a complex architectural challenge.

PURCHASE CONSIDERATIONS

DDN utilizes a transparent capacity-based subscription model (TB/year) that includes all software features, such as encryption and multitenancy, without hidden feature tax add-ons. This all-in licensing simplifies cost predictability for large-scale deployments. The solution is highly flexible, deployable as software on qualified commodity x86 servers, as a turnkey appliance, or as a cloud-native instance, accommodating various buyer preferences from CapEx-heavy on-prem builds to OpEx-centric cloud operations.

The primary buyer profile includes large enterprises and CSPs establishing AI factories that require extreme performance and concurrency. While deployment is automated via Kubernetes operators, the system's advanced engineering focus implies that customers will benefit from DDN's specialized white glove support, particularly for tuning complex parallel I/O paths. Migration is facilitated by nondisruptive live data plane operations, though legacy file protocol support is still maturing compared to unified storage peers.

USE CASES

DDN Infinia supports verticals requiring massive throughput and low latency, such as generative AI development, autonomous driving, and life sciences. Specific use cases include high-performance RAG (retrieval-augmented generation) pipelines where sub-millisecond latency is critical for inference, and multitenant cloud storage services

for which strict logical isolation between competing clients is mandatory. The platform is specifically designed to bridge the gap between object storage scalability and the performance requirements of modern GPUs.

Dell Technologies: Dell ObjectScale

SOLUTION OVERVIEW

Dell ObjectScale is a software-defined object storage platform that now shares a unified codebase with the legacy ECS solution (v4.0). This convergence allows for a consistent S3-compatible software stack deployable across turnkey appliances (such as the all-flash ObjectScale XF960 and HDD-based ObjectScale X560) and ObjectScale Software Defined on validated PowerEdge servers. The architecture supports a global namespace scalable to exabytes, enabling organizations to manage diverse data types from long-term archives to high-performance AI/analytics, across edge, core, and hybrid cloud environments without creating silos.

Dell Technologies prioritizes enterprise-grade stability, security, and proven scale, maintaining a consistent release cadence focused on high-performance AI and cyber resiliency. The solution is designed for highly regulated enterprises, offering a robust zero trust security framework and extensive compliance certifications. Rather than pursuing aggressive, experimental functionality, its roadmap focuses on enhancing cyber resiliency and multicloud tiering. This approach provides a predictable, risk-averse path for modernizing data lakes, allowing customers to support low-latency AI and analytics workloads while maintaining strict governance standards.

Dell Technologies is positioned as a Leader and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Dell Technologies scored well on a number of decision criteria, including:

- **Kubernetes support:** ObjectScale maintains high-tier integration with container orchestration platforms through its support of the emerging COSI standard. This allows for the management of buckets as first-class Kubernetes resources, a key differentiator for cloud-native workflows. While the platform transitioned to a unified, non-Kubernetes-native codebase for improved serviceability in version 4.0, it continues to provide mature CSI drivers and Container Storage Modules (CSM) that automate provisioning and snapshots via standard kubectl commands.

- **Auditing:** ObjectScale provides a native management console that delivers real-time visibility into cluster health, capacity utilization, and performance metrics. The platform features native Grafana integration, allowing administrators to leverage standardized dashboards for detailed telemetry analysis or export monitoring data to external enterprise observability tools. This approach ensures a functional and extensible monitoring foundation, providing IT teams with the necessary insights to manage storage growth and system status effectively.
- **Reporting and analytics:** The solution integrates with CloudIQ to provide AIOps-driven observability across the storage estate. This system analyzes telemetry data to generate predictive capacity forecasts, detect performance anomalies based on learned baselines, and aggregate health checks into a unified score for proactive management.

OPPORTUNITIES

Dell Technologies has room for improvement in a few decision criteria, including:

- **Workload optimization:** The solution utilizes specialized hardware profiles, such as the all-flash XF960 and high-performance software-defined configurations on validated PowerEdge servers, to meet divergent performance needs. While version 4.2 introduces S3 over RDMA to achieve peak throughput, achieving these metrics still requires specific high-end NVMe and networking configurations (RoCE v2). The ability to reach these performance levels on generic, nonvalidated commodity hardware remains a limitation for teams without standardized Dell hardware stacks.
- **Versioning:** ObjectScale supports full S3-compatible versioning, which serves as the technical foundation for data immutability and recovery.The platform utilizes automated, policy-based lifecycle management to prune noncurrent versions and prevent storage bloat. While the system includes default safeguards to manage versioning growth, it relies on administrator-defined policies to align storage consumption with specific business use cases rather than employing fully autonomous internal pruning mechanisms.
- **Storage optimization:** The platform employs flexible erasure coding schemes and selectable inline compression algorithms to maximize density. For multisite deployments (three or more locations), ObjectScale utilizes a global storage efficiency feature based on XOR operations to reduce geo-replication overhead. However, the solution still lacks a global variable-block inline deduplication engine for general S3 objects. This remains a limitation for

customers comparing ObjectScale to dedicated backup appliances that achieve significantly higher data reduction ratios through deduplication.

PURCHASE CONSIDERATIONS

Licensing for ObjectScale is based on capacity (per TB) and node counts, offering a predictable cost structure that avoids the variable egress fees and API request charges common in public cloud models. The availability of software-defined options allows enterprises to leverage existing hardware supply chains, further optimizing TCO.

The solution targets large enterprises, CSPs, and federal agencies rather than SMBs. It supports a flexible deployment matrix, including turnkey appliances for organizations desiring a black box support model and software-defined configurations for engineering-led teams.

Deployment complexity varies by model. While appliances remain a turnkey black box solution, the software-defined model now supports bare metal deployment, removing the previous requirement for specialized Kubernetes expertise to maintain the underlying infrastructure. This architectural shift significantly simplifies the operational profile for organizations leveraging existing hardware supply chains. Migration is a notable strength. The unified codebase allows existing elastic cloud storage (ECS) customers to upgrade to ObjectScale via a nondisruptive data-in-place process, eliminating the migration tax often associated with platform modernization.

USE CASES

Dell Technologies supports verticals requiring extreme scale and compliance, including finance, federal government, and global 2000 enterprises. Specific use cases include high-performance AI model training, enabled by S3 over RDMA and flash appliances to feed GPU clusters; massive data lakes for analytics; and long-term archival retention. The solution is engineered to serve as a unified repository for both primary high-speed data and deep archival content within a single namespace.

DigitalOcean: Spaces*

SOLUTION OVERVIEW

DigitalOcean simplifies cloud infrastructure for developers, startups, and SMBs, focusing on reducing complexity and providing predictable costs. The vendor has maintained a consistent trajectory of expanding its platform capabilities to serve cloud-native applications

without the cognitive overhead associated with hyperscale providers. DigitalOcean recently enhanced its storage portfolio with the introduction of a cold storage tier, signaling a move toward broader data lifecycle management.

Spaces is the vendor's S3-compatible object storage solution, delivered exclusively as a managed service. It is not a standalone product but a key component of the broader DigitalOcean platform, designed to integrate with Droplets (compute), App Platform (PaaS), and DigitalOcean Kubernetes. The offering includes a built-in content delivery network (CDN) at no additional base cost, tightly coupling storage with delivery. Spaces focuses on delivering core object storage functionality of storage, retrieval, and serving with high reliability and minimal configuration.

The solution will look and feel different over the contract lifecycle. DigitalOcean delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. Primary emphasis on emerging features and rapid development to address gaps ensures the platform evolves alongside modern developer requirements.

DigitalOcean is positioned as an Entrant and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

DigitalOcean scored well on a number of decision criteria, including:

- **Workload optimization:** Spaces combines automated lifecycle policies with a native CDN to optimize both storage costs and delivery performance. Users can configure rules to automatically transition older objects to the cold storage tier, reducing costs for infrequently accessed data. Concurrently, the integrated CDN offloads traffic from the origin bucket to edge nodes, accelerating content delivery for web-scale workloads without complex external configurations.
- **Versioning:** The platform supports full S3-compatible object versioning, allowing users to preserve, retrieve, and restore every version of every object stored. This capability provides essential protection against accidental deletion and application failures. Users can list versions and implement rollback strategies using standard S3 tools, ensuring data integrity at the application level.
- **Public cloud integration:** Spaces leverages high-fidelity S3 API compatibility to function effectively within multicloud architectures. While it does not offer native bidirectional sync tools, the API support allows seamless integration with standard data migration and backup tools like AWS DataSync, rclone, and third-party

backup solutions. This interoperability makes it a viable tier for offloading data from other public clouds or acting as a secondary backup target.

OPPORTUNITIES

DigitalOcean has room for improvement in a few decision criteria, including:

- **Ransomware protection:** The solution currently lacks support for S3 Object Lock (WORM), a critical feature for immutable data protection. While versioning allows for recovery from accidental deletion, the absence of governance or compliance modes means data remains vulnerable to deletion by compromised administrator credentials. Organizations with strict regulatory requirements for immutable storage will find this limiting.
- **Reporting and analytics:** Native observability tools are restricted to basic dashboard graphs showing bandwidth usage and request counts. The platform does not currently offer deep storage analytics, such as heatmaps, prefix-level usage breakdowns, or storage lens capabilities. Users seeking granular insights into data access patterns or predictive capacity planning must rely on external monitoring solutions.
- **Kubernetes support:** Integration with Kubernetes is primarily focused on block storage via the CSI driver. There is no official COSI driver or native CSI support specifically for provisioning Spaces buckets. Users needing to mount object storage as persistent volumes must rely on community-maintained adapters or Filesystem in Userspace (FUSE)-based workarounds, which lack the enterprise support guarantees of the block storage integration.

PURCHASE CONSIDERATIONS

The pricing model for Spaces is notably transparent and predictable, distinguishing it from competitors with complex fee structures. The base subscription includes storage and a generous outbound transfer allowance, with no charges for standard API requests (GET, PUT, LIST). This elimination of transaction fees makes cost forecasting exceptionally simple, particularly for read-heavy or chatty applications that would incur significant costs on hyperscale platforms. The absence of complex tiered egress fees further simplifies the TCO calculation.

Spaces is ideally suited for SMBs, developers, and independent software vendors (ISVs) who prioritize ease of use and price predictability over infinite configurability. It is less suited for large enterprises requiring complex identity and access management (IAM)

hierarchies or strict regulatory compliance features like WORM. Deployment is instant via the cloud console, requiring no infrastructure provisioning. Migration is straightforward for S3-native applications, as most codebases require only an endpoint configuration change to switch providers. Professional services are generally not required due to the platform's batteries included design philosophy.

USE CASES

DigitalOcean primarily supports cloud-native developers, startups, and media-centric businesses. Key use cases include hosting static website assets, serving video and media content via the integrated CDN, and acting as a backend for data-intensive applications. The platform is also widely used by MSPs for backup repositories and media management solutions because the predictable pricing model protects margins. Its architecture supports these workloads by combining low-friction storage with high-performance edge delivery.

Everpure: FlashBlade

SOLUTION OVERVIEW

Everpure (a recent rebranding from Pure Storage) focuses on delivering the all-flash data center, aiming to retire mechanical disk through high-performance, solid-state architectures. The primary object storage offering is the FlashBlade portfolio, comprising the performance-optimized FlashBlade//S and the capacity-optimized FlashBlade//E. These systems use a Unified Fast File and Object (UFFO) strategy, consolidating workloads onto a single platform. For containerized environments, the vendor integrates Portworx, a software-defined data services platform.

The solution addresses diverse needs from AI training and high-frequency analytics to rapid recovery backup targets. Everpure creates a distinct value proposition by eschewing traditional tradeoffs between speed and capacity, leveraging DirectFlash technology to manage NAND flash memory directly without SSD controllers. This architecture supports the Elastic Core, a distributed metadata engine designed to handle massive concurrency without the bottlenecks typical of gateway-based object stores.

The solution is likely to look and feel materially different over the contract lifecycle, as Everpure delivers an aggressive roadmap that may invite disruption. The vendor's rapid adoption of emerging capabilities—including quad-level cell (QLC) flash for capacity-tier displacement, S3-over-RDMA, and GPUDirect Storage integration—combined with the strategic acquisition of Portworx to extend into Kubernetes data services, reflects a pattern of M&A and accelerated feature development designed to keep the platform ahead of evolving AI and cloud-native requirements.

Everpure is positioned as a Leader and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Everpure scored well on a number of decision criteria, including:

- **Ransomware protection:** The platform implements SafeMode, which creates immutable snapshots that cannot be deleted or modified, even by administrators with root privileges, until a retention timer expires. This is complemented by native support for S3 Object Lock in both compliance and governance modes, validated for use with major backup vendors to provide a robust defense against encryption attacks.
- **Reporting and analytics:** The Pure1 management plane offers AI-driven insights that extend beyond basic monitoring. It provides capacity forecasting, workload simulation, and fingerprint issue detection that analyzes telemetry across the global fleet to predict and prevent disruptions. This creates a zero touch operational model whereby storage management is largely predictive rather than reactive.
- **Kubernetes support:** Through the integration of Portworx, the solution offers a mature data services layer for containerized applications. It supports the emerging COSI standard, enabling developers to provision buckets via Kubernetes manifests. Additionally, topology-aware scheduling ensures compute pods are placed near their data to minimize latency in hybrid environments.

OPPORTUNITIES

Everpure has room for improvement in a few decision criteria, including:

- **Auditing:** While the system supports standard logging of system events and S3 access logs exportable to syslog, it lacks deep, native introspection capabilities. Detailed analysis of user behavior or granular audit trails often requires integration with external SIEM tools rather than being surfaced directly within the native administrative interface.
- **Storage optimization:** The architecture uses compression and pattern removal but explicitly omits global deduplication for object storage workloads. This design choice prioritizes restore throughput and read performance but results in lower storage efficiency ratios compared to deduplication-centric backup appliances, potentially affecting the TCO for highly repetitive datasets.

- **Versioning:** The platform supports S3 bucket versioning to manage object history and recovery, but the lifecycle management for these versions relies on external dependencies for advanced tiering. Moving noncurrent versions to lower-cost external cloud tiers typically requires third-party software (such as Komprise) rather than being handled by a comprehensive, native policy engine within the operating system.

PURCHASE CONSIDERATIONS

The Everpure commercial model is established for its transparency and flexibility. The Evergreen//One subscription transforms storage procurement into a utility service, for which costs are aligned with effective usage rather than raw hardware capacity. This model includes SLAs for performance and uptime, and it eliminates the traditional forklift upgrade cycle by including nondisruptive hardware refreshes in the subscription.

The solution is primarily targeted at large enterprise and service provider markets. The appliance-based delivery model creates a high entry point in terms of capital and capacity, making it generally unsuitable for SMBs or edge deployments requiring small-scale commodity hardware.

Deployment is simplified through the appliance form factor, which arrives as a tightly integrated unit, reducing the complexity often associated with SDS on bare metal. However, this hardware dependency means the solution cannot be deployed as software-only on existing third-party servers (outside of the Portworx integration for Kubernetes), distinguishing it from fully software-defined competitors.

USE CASES

Everpure supports high-performance verticals such as media and entertainment, life sciences, and financial services. The platform is specifically optimized for AI and ML workflows, wherein features like S3-over-RDMA and NVIDIA GPUDirect Storage eliminate CPU bottlenecks to keep GPUs saturated. Additionally, the system is widely used for rapid restore scenarios in modern data protection, where the lack of deduplication rehydration allows for recovery speeds that significantly exceed those of traditional backup appliances. The architecture also serves as a backend for Splunk SmartStore, supporting high-speed analytics queries.

Hitachi Vantara: Virtual Storage Platform One Object

SOLUTION OVERVIEW

Hitachi Vantara, a subsidiary of Hitachi, Ltd., is a prominent player in the enterprise infrastructure market, pivoting from a product-centric hardware focus to a unified platform strategy. Known for mission-critical reliability, the vendor has integrated its storage portfolio under the VSP One banner, dissolving silos separating block, file, and object storage. Recent strategic moves include the transition from the legacy appliance-based Hitachi Content Platform (HCP) to the software-defined, microservices-based Virtual Storage Platform One (VSP One) Object.

The primary solution, VSP One Object, is a scalable, software-defined object store available as a standalone solution or part of the broader VSP One data platform. Relevant product SKUs include VSP One SDS, VSP One Block backends, and the established Hitachi Content Platform (HCP) portfolio, including HCP Anywhere Enterprise, for extended edge and compliance use cases. Hitachi Vantara employs a Platform Play strategy, aiming to provide a holistic data foundation for hybrid cloud, analytics, and AI workloads rather than a niche point solution.

The solution will look and feel different over the contract lifecycle. Hitachi Vantara delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. Primary emphasis on emerging features like native S3 Tables, agentic AI support, and advanced metadata services drives this innovation focus.

Hitachi Vantara is positioned as a Leader and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Hitachi Vantara scored well on a number of decision criteria, including:

- **Ransomware protection:** The solution delivers a multilayered defense strategy centered on S3 Object Lock for immutability. It augments this with HCP Anywhere Enterprise Ransom Protect, a feature using ML algorithms to monitor data access patterns and detect behavioral anomalies such as mass encryption events in real time to trigger automated user lockdowns.
- **Storage optimization:** Hitachi Object Storage combines efficient erasure coding (for example, 20+6 schemes) with granular, bucket-level deduplication and compression. Uniquely, when paired with VSP One Block backends, it leverages hardware-offloaded advanced data reduction (ADR) to deliver a 4:1 data reduction guarantee without compromising performance.

- **Public cloud integration:** The platform supports tiering to the cloud, allowing for the movement of data from on-prem Hitachi object storage and public cloud targets, such as AWS, Azure, and Google Cloud Platform (GCP). It can also be deployed as a software-defined instance within the public cloud for consistent disaster recovery or hybrid workflows.

OPPORTUNITIES

Hitachi Vantara has room for improvement in a few decision criteria, including:

- **Workload optimization:** While the solution supports intelligent tiering and high-performance analytics via S3 Tables, it currently lacks direct GPU data paths. The roadmap includes S3 Vectors and GPUDirectStorage (GDS) support, but presently, AI training data must traverse the standard CPU stack, potentially introducing latency compared to specialized HPC storage.
- **Versioning:** The platform offers robust S3-compatible versioning integrated with lifecycle management and Object Lock compliance. However, while version-aware restoration is precise, the system could further enhance user experience by simplifying the management of massive version histories at exabyte scale, which currently relies on standard S3 list/delete operations rather than more advanced automated pruning policies found in some niche competitors.
- **Kubernetes support:** VSP One Object runs on a Kubernetes distribution, but its external-facing support for developers is limited. While it supports Ansible, Swagger, and AWS SDKs, it lacks a mature Operator for dynamic bucket provisioning. The solution currently relies on CSI drivers focused on block/file storage and lacks a mature operator or CRDs for dynamic, declarative object bucket provisioning, which is a roadmap item for 2026.

PURCHASE CONSIDERATIONS

Hitachi Vantara's licensing model prioritizes transparency, using a capacity-based approach by which fees are tied to consumed usable capacity rather than raw hardware or compute nodes. This prevents incursion of the capacity tax often associated with performance upgrades. The EverFlex consumption program offers flexible, cloud-like pricing (XaaS) to further align costs with usage.

The solution is well suited for large enterprise buyers, particularly in regulated industries like finance and government, given its compliance heritage. Deployment is flexible, offered as SDS on a VM commodity hardware or integrated appliances, though it typically

requires a displacement or greenfield implementation rather than a simple module add-on. Migration is facilitated by the platform's S3 compatibility, but moving from legacy proprietary archives may require professional services, which Hitachi Vantara supports with a strong partner ecosystem.

USE CASES

Hitachi Vantara supports vertical-specific needs in highly regulated sectors such as finance (SEC/FINRA compliance), healthcare (imaging archives), and government (FIPS 140-2 security) through robust auditing and immutability features. Primary use cases include on-prem data lakehouse, AI/ML, and analytics leveraging native S3 Tables and Apache Iceberg integration alongside active archiving for hybrid cloud environments. This correlates with its Platform Play designation. It serves as a unified foundation for diverse enterprise workloads rather than a single-purpose repository.

IBM: IBM Cloud Object Storage (COS), IBM Storage Ceph*

SOLUTION OVERVIEW

IBM offers a dual-platform object storage portfolio comprising IBM Cloud Object Storage (COS) and IBM Storage Ceph. COS targets exabyte-scale archival workloads using Information Dispersal Algorithms, while Ceph addresses high-performance Kubernetes-native requirements through a software-defined architecture. These storage engines are integrated within the IBM Storage Fusion layer for orchestration and secured by the IBM Storage Defender suite. The portfolio also includes Red Hat OpenShift Data Foundation (ODF), built on Ceph, to support container-native storage needs. IBM leverages its public cloud to offer COS as a SaaS, while IBM Cloud Satellite enables deployment of this SaaS on-prem or at the edge. The strategy emphasizes convergence through common management and data services rather than unifying the underlying storage engines.

The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. IBM values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement.

IBM is positioned as a Leader and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

IBM scored well on a number of decision criteria, including:

- **Kubernetes support:** IBM utilizes the Rook operator to manage IBM Storage Ceph within Kubernetes environments. Unlike standard CSI drivers, Rook automates deployment, scaling, upgrades,

and disaster recovery, effectively treating the storage cluster as a managed Kubernetes application. Additionally, IBM Storage Fusion enables application-aware backups by snapshotting entire namespaces, including persistent volume claims (PVCs), to facilitate mobility between on-prem and cloud clusters.

- **Ransomware protection:** The solution combines hardware, software, and logical protections. FlashCore Modules (FCM) perform inline entropy analysis to detect encryption activity at the block level in real time. This hardware data feeds IBM Storage Defender, which can trigger responses like immutable snapshots or bucket lockdowns. These features operate alongside native S3 Object Lock in compliance mode to ensure data immutability.
- **Storage optimization:** IBM COS employs a proprietary Information Dispersal Algorithm (IDA) that slices and distributes data across nodes, achieving high durability with approximately 1.7x storage overhead compared to the 3x often required for replication. IBM Storage Ceph utilizes the BlueStore backend to support inline compression algorithms like LZ4 and Snappy, and integrates with IBM Storage Protect for source-side deduplication.

OPPORTUNITIES IBM has room for improvement in a few decision criteria, including:

- **Auditing:** The platform generates detailed audit logs for management and data events, supporting real-time streaming to external SIEM tools like Splunk and QRadar. However, the core auditing service functions primarily as a capture and export engine. Advanced capabilities such as predictive issue prevention or AI-driven anomaly analysis are not intrinsic to the audit log parser itself and must be integrated with the separate IBM Storage Defender or Storage Insights platforms.
- **Versioning:** IBM supports S3-compliant versioning to retain object variants and integrates these with lifecycle policies to manage storage costs. While the functionality is comprehensive, the existence of two distinct object platforms, COS and Ceph, means administrators must manage versioning policies through different interfaces. This separation creates operational divergence compared to single-stack architectures, requiring distinct configuration workflows for each environment.
- **Reporting and analytics:** IBM Storage Insights Pro offers a centralized dashboard for health monitoring, capacity planning, and automated support ticketing. Despite this unified view, granular troubleshooting often necessitates accessing the specific element managers, such as the IBM COS Manager or Ceph Dashboard.

While efforts are underway to consolidate management within Fusion, the current requirement to toggle between the high-level analytics tool and platform-specific interfaces creates a disjointed administrative experience.

PURCHASE CONSIDERATIONS

IBM offers transparent pricing models, specifically the Smart Tier for cloud storage, which automates classification based on activity to eliminate retrieval fees. For on-prem software, the IBM Storage Suite provides capacity-based licensing per TB that includes all features, avoiding complex module-specific costs. The solution targets large enterprise and service provider markets, as the minimum node requirements for efficient IDA in COS and the operational skills needed for Ceph generally preclude SMB adoption. Deployment options are versatile, supporting turnkey physical appliances, software-only subscriptions, and fully managed SaaS via IBM Cloud Satellite. This range enables organizations to deploy the same software stack across CapEx and OpEx models.

USE CASES

IBM supports regulated verticals such as banking, healthcare, and government. The portfolio addresses two primary use cases: exabyte-scale active archiving, by which IBM COS serves as a repository for backups and compliance records; and the data lakehouse, by which IBM Storage Ceph provides high-performance backend storage for AI and analytics workloads utilizing watsonx. data. This structure allows the vendor to manage the data lifecycle from high-speed ingestion to long-term retention.

Infinidat: InfiniBox

SOLUTION OVERVIEW

Infinidat specializes in enterprise-class storage solutions designed for multi-petabyte scale, high availability, and economic efficiency. Rather than relying solely on all-flash architectures, the vendor uses a software-defined approach centered on its Neural Cache architecture, which leverages deep learning algorithms to optimize performance across high-density hardware. The entire portfolio, including InfiniBox and the all-flash InfiniBox SSA, runs on the unified InfuzeOS. This architecture delivers native, concurrent block, file, and S3-compatible object storage services from a single platform without the use of external gateways. By avoiding translation layers, Infinidat provides a consistent management experience and optimizes performance across all supported protocols. The portfolio also includes InfiniGuard for purpose-built data protection.

The InfiniBox is a consolidated platform designed to eliminate storage silos. By integrating native S3 object storage alongside Fibre Channel and NFS protocols, Infinidat targets large enterprises and service providers seeking to centralize diverse workloads on a unified infrastructure. The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. Infinidat values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement, evidenced by its reliance on the battle-tested InfiniRAID architecture and 100% availability guarantees.

Infinidat is positioned as a Leader and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Infinidat scored well on a number of decision criteria, including:

- **Auditing:** InfuzeOS provides granular visibility into system access and data operations, a critical requirement for multitenant service provider environments. The system captures detailed event logs that can be streamed in real time to external Syslog servers or SIEM platforms like Splunk. This capability allows security operations teams to correlate storage access patterns with broader threat intelligence, facilitating the detection of anomalies such as mass deletions or unauthorized encryption events.
- **Versioning:** The platform supports a robust S3-compatible object versioning implementation that goes beyond simple recovery. It integrates tightly with lifecycle management policies, allowing administrators to define automated expiration or transition rules for noncurrent object versions. This integration supports strict compliance and data protection workflows, ensuring previous iterations of data remain retrievable for regulatory discovery or operational rollback without manual intervention.
- **Storage optimization:** Efficiency is a core architectural tenet, achieved through a combination of high-capacity hard disk drives and the patented Neural Cache. The system employs always-on inline data reduction, including compression, which operates without performance penalty due to the DRAM-centric I/O path. The underlying InfiniRAID layout maximizes usable capacity, scaling to millions of gigabytes per rack, while minimizing the overhead typically associated with dual-parity protection schemes.

OPPORTUNITIES

Infinidat has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While the vendor offers a mature, Red Hat-certified CSI driver for block and file workloads, it currently lacks support for the COSI standard. This limits the ability of cloud-native developers to provision object buckets dynamically via Kubernetes APIs. Organizations require administrative intervention to create buckets via the UI or REST API before they can be consumed by containerized applications, contrasting with the self-service models of software-only object stores.
- **Ransomware protection:** InfiniSafe provides exceptional capabilities, including immutable snapshots and forensic scanning, but these features often require a dedicated or fenced forensic environment for safe recovery and analysis. While powerful, this approach necessitates careful architectural planning to ensure the isolated environment is properly resourced and configured. Organizations must actively manage these forensic reserves to ensure rapid recovery mechanisms function as intended during a live cyber event.
- **Reporting and analytics:** The InfiniVerse platform delivers deep, AI-driven insights into capacity planning and performance forecasting, yet it operates primarily as a cloud-based AIOps service. While this aggregates telemetry effectively for global oversight, organizations with strict air-gapping requirements may face limitations in accessing the full suite of predictive analytics. Ensuring equivalent granular visibility for dark site deployments remains an area where on-prem tooling could be expanded to match the SaaS experience.

PURCHASE CONSIDERATIONS

Infinidat is notable for its high degree of cost transparency and simplified licensing structure. Unlike competitors that charge separately for advanced features, the InfiniBox pricing is all-inclusive, bundling replication, snapshots, and cyber resilience capabilities into the base cost. The vendor offers flexible consumption models, including capacity on demand and FLX, which allow enterprises to align expenses with actual usage, effectively shifting storage from a capital expenditure to an operational expense.

The solution is engineered specifically for large enterprises, CSPs, and MSPs. The physical infrastructure requirements (typically a full 42U rack scale) and the multi-petabyte entry point make it unsuitable for SMBs or edge locations requiring small-footprint storage. Deployment involves the delivery of a preintegrated appliance, simplifying the rack and stack phase but lacking the software-only flexibility found in commodity server-based object stores. Migration is facilitated by the vendor's White Glove service, which assists with the transition of massive datasets from legacy arrays.

USE CASES

Infinidat primarily supports data-intensive verticals such as finance, telecommunications, and healthcare, for which reliability cannot be compromised. The platform excels in consolidation use cases, allowing a bank, for example, to host high-performance databases, general-purpose file shares, and S3-based backup repositories on a single array. Additionally, the Neural Cache architecture makes it highly effective for AI and analytics workflows that require low-latency access to massive datasets, while the InfuzeOS Cloud Edition extends these capabilities to hybrid cloud architectures for disaster recovery.

MinIO: AIStor

SOLUTION OVERVIEW

MinIO is a high-performance, software-defined object storage vendor headquartered in Redwood City, California. The company focuses on serving as the primary data persistence layer for AI, ML, and modern data lakehouse architectures, prioritizing throughput and parallelism over legacy protocol support. MinIO AIStor is a unified, software-defined object storage platform delivered as a single binary, distinct from portfolio suites that require integrating separate products.

MinIO delivers a comprehensive AI data infrastructure that orchestrates data flow from edge to core to cloud. The solution continues to push the boundaries of what object storage is expected to do, consistently introducing capabilities such as S3 over RDMA, the PromptObject API, and native Iceberg catalog support that anticipate the requirements of AI-driven data estates. MinIO's willingness to deprecate its own features in service of architectural modernization reflects a vendor that prioritizes engineering velocity over incremental refinement.

MinIO is positioned as a Leader and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

MinIO scored well on a number of decision criteria, including:

- **Kubernetes support:** MinIO utilizes a Kubernetes Operator that manages tenant creation, expansion, and upgrades through declarative CRDs, aligning with GitOps workflows. To address the latency often found in standard CSI drivers, MinIO utilizes DirectPV, a specialized driver that manages local drives directly to ensure the object store operates at the speed of local media.
- **Workload optimization:** The platform implements granular QoS controls, allowing administrators to set bandwidth and IOPS limits

at the bucket or prefix level to prevent noisy neighbor issues. MinIO supports the AWS S3 Express One Zone API, which employs directory buckets and atomic append operations to support high-frequency transaction logs and low-latency workloads.

- **Versioning:** MinIO treats every overwrite or delete operation as a new object version identified by a unique ID, providing an infinite history of the object state for point-in-time recovery. This capability integrates tightly with Object Lock and Lifecycle Management, allowing users to define policies that automatically transition or expire noncurrent versions to manage capacity.

OPPORTUNITIES

MinIO has room for improvement in a few decision criteria, including:

- **Auditing:** The solution captures granular events for every API interaction and streams them in real time to external targets like Kafka and Elasticsearch. However, MinIO does not provide built-in audit log search capabilities within the console itself, relying instead on these external systems to perform search, analysis, and visualization of the audit streams.
- **Ransomware protection:** MinIO employs a defense-in-depth architecture featuring S3 Object Lock for immutability and active anomaly detection that monitors for spikes in deletion requests. While the system integrates eBPF to block deletion calls at the kernel level, the vendor notes continued work on enhancing predictive analytics capabilities to further strengthen end-to-end deployment security and detection.
- **Reporting and analytics:** The MinIO Console provides a centralized view of cluster health and throughput, exposing exhaustive metrics via a Prometheus V3 endpoint for external visualization. The vendor acknowledges room to improve by adding AI telemetry capabilities to further simplify Day 2 operations and enhance native predictive analytics for capacity and hardware health.

PURCHASE CONSIDERATIONS

MinIO employs a transparent, capacity-based subscription model priced per tebibyte per month, which bundles the software license with the SUBNET support subscription. This all-inclusive approach avoids complex tiered add-ons or per-seat licensing, allowing for precise TCO modeling. The solution is viable for a wide range of buyers, from small-to-medium businesses to large enterprises and CSPs.

Deployment is notably flexible. As a software-defined and hardware-agnostic solution, AIStor runs consistently across bare metal, virtual

machines, containers, and public cloud instances. It is delivered as a single static binary of roughly 100MB, reducing dependency complexity compared to modular legacy suites. Support is delivered through the SUBNET portal, which provides direct-to-engineer communication and automated fleet health checks rather than a traditional tiered support hierarchy.

USE CASES

MinIO supports verticals including autonomous vehicles, healthcare, manufacturing, and logistics, handling massive volumes of unstructured data for these sectors. Primary use cases center on AI/ML workflows, specifically LLM training and RAG pipelines, and modern data lakehouses using the Apache Iceberg format. The platform also supports migration from the Hadoop Distributed File System (HDFS) and enterprise data warehouses, consolidating data across environments. It serves as a comprehensive foundation for diverse high-performance data workloads.

NetApp: NetApp StorageGRID

SOLUTION OVERVIEW

NetApp is a data management authority focused on unifying storage infrastructure across on-prem and hybrid cloud environments. Its StorageGRID solution is a software-defined object storage platform designed for stability, massive scalability, and complex policy management rather than raw speed alone. StorageGRID enables organizations to manage unstructured data as a comprehensive fabric, integrating deeply with the broader NetApp ecosystem, including ONTAP and the BlueXP control plane. The solution creates a single global namespace that can span up to 16 geographically dispersed sites, abstracting physical locations for applications and users.

StorageGRID is available as purpose-built hardware appliances (such as the all-flash SGF6112 and hybrid SG6000 series), as virtual appliances on VMware, or as software-only containers for bare metal deployments. The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. NetApp values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement, targeting regulated industries and large enterprises requiring long-term data retention with active accessibility.

NetApp is positioned as a Leader and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

NetApp scored well on a number of decision criteria, including:

- **Auditing:** StorageGRID generates comprehensive audit logs for every object transaction (GET, PUT, DELETE) and management action, satisfying stringent compliance requirements in regulated sectors. These logs are not locked in proprietary databases but are accessible and exportable. The solution integrates directly with analytics platforms like Splunk, enabling security operations centers to visualize access patterns and detect anomalies such as sudden spikes in data egress or unauthorized access attempts.
- **Public cloud integration:** The platform supports deep bidirectional integration with major public clouds, enabling true hybrid workflows. Features such as Cloud Storage Pools allow the system to tier objects to Amazon S3, Azure Blob, or Google Cloud Storage based on information lifecycle management (ILM) policies. Additionally, the CloudMirror service replicates objects to external S3 buckets for disaster recovery or content distribution, allowing data to be processed by cloud-native serverless functions while the master copy remains on-prem.
- **Versioning:** StorageGRID supports full S3-compatible bucket versioning, allowing users to preserve, retrieve, and restore every version of every object stored. The system integrates versioning with its sophisticated ILM policy engine, enabling granular rules for version retention. For example, administrators can configure policies to keep current versions on high-performance flash nodes while automatically moving noncurrent versions to capacity-optimized archives or deleting them after a set period.

OPPORTUNITIES

NetApp has room for improvement in a few decision criteria, including:

- **Storage optimization:** StorageGRID delivers reliable compression and highly configurable erasure coding schemes that can protect data with as little as 20–30% overhead. However, the platform lacks native global deduplication, which is the ability to identify and eliminate redundant data blocks across the entire namespace. NetApp's architectural position is that deduplication is best performed at the source, such as by the backup application or via ONTAP FabricPool tiering, before data reaches the object store. While this approach is architecturally valid, it means that for primary workloads ingesting unreduced data, StorageGRID may consume more raw capacity than deduplication-native competitors.
- **Workload optimization:** StorageGRID features object-native QoS and intelligent data placement, but it lacks the direct kernel-bypass

mechanisms found in specialized HPC solutions. Specifically, while NetApp supports NVIDIA GPUDirect Storage on its ONTAP NAS systems, this capability has not bridged to the StorageGRID S3 architecture. Consequently, data must traverse the standard CPU and memory path, introducing latency that distinguishes it from specialized parallel file systems designed for AI training.

- **Reporting and analytics:** The platform provides exceptional visualization via built-in Prometheus exporters and Grafana dashboards, but active anomaly detection relies on external portfolio integration. Real-time detection of mass deletions or encryption patterns, which is critical for ransomware defense, is primarily handled through integration with NetApp Cloud Insights (Data Infrastructure Insights) rather than being a native, self-contained capability of the grid software itself.

PURCHASE CONSIDERATIONS

NetApp employs a transparent and flexible licensing strategy centered on the NetApp Keystone model. This allows customers to consume StorageGRID capacity as a service, paying a per-TiB rate for capacity used with the ability to burst and scale down, effectively shifting CapEx to OpEx while retaining data sovereign control. Traditional perpetual licensing is also available, and the decoupling of software from hardware allows for the mixing of different appliance generations and models within a single grid.

The solution is engineered for scale, typically starting at hundreds of terabytes or petabytes, making the architectural complexity and minimum node requirements generally prohibitive for SMBs. It is explicitly targeted at large enterprises, CSPs, and MSPs. Deployment is flexible, supporting physical appliances, virtual machines, and software-only docker containers. However, the architecture's reliance on a distributed Cassandra database for metadata means scaling requires careful planning of both storage and compute nodes. Migration from legacy systems is facilitated by the global namespace and nondisruptive hardware refresh capabilities.

USE CASES

NetApp StorageGRID serves the large enterprise and CSP markets, with a strong foothold in regulated verticals such as healthcare, finance, and government. Its deep auditing and compliance features (SEC 17a-4(f) certification) make it a primary choice for long-term retention of sensitive data. The solution supports active archives and data lakes, particularly where data must be geo-distributed across multiple sites for durability. It also acts as a central repository for hybrid cloud workflows, tiering cold data from high-performance ONTAP arrays via FabricPool.

Nutanix: Nutanix Unified Storage

SOLUTION OVERVIEW

Nutanix is an enterprise cloud platform provider that has evolved from hyperconverged infrastructure (HCI) to offer a broader portfolio of hybrid cloud data services. Nutanix Unified Storage (NUS) is a software-defined data management platform that aggregates block, file, and object storage services onto a single distributed fabric. The object storage component, Nutanix Objects, is deployed as a containerized microservice atop the Nutanix Cloud Platform (NCP), separating S3 service interfaces from the underlying persistence layer to allow independent scaling.

The solution enables organizations to consolidate diverse workloads from backup archives to high-performance analytics onto a unified hardware footprint. By leveraging the Nutanix Distributed Storage Fabric (DSF), NUS provides consistent data protection and management across on-prem clusters and public cloud deployments via Nutanix Cloud Clusters (NC2), supporting a federated global namespace and automated tiering to major public cloud providers. The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. Nutanix values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement.

Nutanix is positioned as a Leader and Outperformer in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Nutanix scored well on a number of decision criteria, including:

- **Workload optimization:** Nutanix Objects is architected to support high-throughput workloads beyond archival storage, utilizing intelligent metadata placement on NVMe tiers. In MLPerf Storage v1.0 benchmarking, Nutanix Unified Storage demonstrated the platform's overall high-throughput capabilities; notably, Nutanix Objects was the first object storage solution to submit verified results in the benchmark, successfully saturating dozens of accelerators per node to validate its suitability for AI training workloads.
- **Auditing:** The solution integrates with Nutanix Data Lens, a SaaS-based management platform that provides visualized auditing and granular visibility into access patterns. This capability allows administrators to track user behavior, identify anomalies, and analyze storage consumption trends, including data age and dark data distribution across global clusters.
- **Versioning:** The platform supports standard S3 versioning to maintain multiple object variants, protecting against accidental

deletion or overwrites. This feature functions in tandem with S3 Object Lock (WORM) capabilities, which are validated for SEC Rule 17a-4(f) compliance, enabling strict governance and compliance mode retention policies for regulated industries.

Nutanix is classified as an Outperformer based on its rapid advancement in key capabilities over the last year. The vendor has significantly enhanced its security posture with granular auditing and ransomware protection while simultaneously delivering high-performance features like AI-optimized data paths and global namespace support, effectively bridging the gap between core enterprise storage and modern cloud-native requirements.

OPPORTUNITIES

Nutanix has room for improvement in a few decision criteria, including:

- **Kubernetes support:** Nutanix provides a production-ready COSI driver for declarative bucket provisioning. However, COSI is an emerging standard relative to the mature CSI, and adoption requires organizations to integrate new CRDs like BucketClaim into their existing DevOps workflows, which may differ from established persistent volume practices.
- **Ransomware protection:** The platform offers active cyber defense capabilities, including automated threat containment and behavioral analysis to block attacks within minutes. While active cyber defense features were previously tied to a SaaS-based model, the introduction of Nutanix Data Lens as an on-prem deployment addresses requirements for air-gapped environments and strict data sovereignty. This expansion allows organizations to maintain real-time behavioral analysis and automated threat containment without transmitting telemetry to a cloud management plane.
- **Storage optimization:** Nutanix Objects utilizes adaptive erasure coding (EC-X) and inline compression to maximize capacity efficiency. However, the vendor explicitly recommends disabling deduplication for object storage workloads to avoid unnecessary metadata overhead, which may result in lower storage efficiency compared to purpose-built backup appliances that rely on global deduplication for massive capacity reduction.

PURCHASE CONSIDERATIONS

Nutanix Unified Storage uses a capacity-based licensing model that is transferable across on-prem and public cloud environments, simplifying cost predictability compared to component-based pricing. The solution is generally targeted at mid-market and large enterprise organizations seeking to reduce infrastructure silos through

hyperconvergence. Buyers should note that while the core software runs on standard server hardware or bare metal cloud instances, the unified architecture replaces traditional dedicated storage arrays, representing a shift in deployment strategy. Management is centralized through the Prism console, designed to lower operational overhead, though the full suite of advanced analytics requires the Data Lens SaaS component.

USE CASES

Nutanix Unified Storage is suitable for regulated verticals like financial services due to its SEC-validated WORM support. The solution supports high-performance AI data lakes, evidenced by MLPerf validation for image segmentation and training workloads. Furthermore, its suitability for modern analytics is validated through integrations with third-party platforms such as Dremio, Snowflake, and Vertica, enabling these engines to query unstructured data directly on Nutanix Objects. Additionally, it serves as a primary immutable target for enterprise backup applications such as Veeam and Commvault, replacing tape or dedicated appliances within a consolidated platform architecture.

OSNexus: QuantaStor

SOLUTION OVERVIEW

OSNexus addresses the complexity of open source storage with QuantaStor, a unified SDS platform that integrates scale-out object (Ceph) and scale-up file/block (ZFS) technologies into a manageable grid. The solution is designed to democratize hyperscale architecture for the enterprise, allowing organizations to deploy massive multi-petabyte private clouds on commodity hardware from vendors like Seagate, Supermicro, and Western Digital. QuantaStor is a comprehensive platform rather than a niche tool, unifying block, file, and object protocols within a single storage grid management plane that federates up to 64 appliances.

Strategically, OSNexus adheres to a Maturity orientation. Rather than racing to implement experimental features, the vendor prioritizes stability, security certifications (FIPS 140-2), and seamless hardware integration. The solution delivers a consistent, secure experience suited for regulated industries, focusing on operationalizing proven technologies like Ceph RADOS rather than reinventing the filesystem.

OSNexus is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

OSNexus scored well on a number of decision criteria, including:

- **Ransomware protection:** QuantaStor delivers exceptional cyber resilience by combining S3 Object Lock (governance and compliance modes) with active defense mechanisms. Beyond standard immutability, the platform features a proactive Anomaly Detection system that monitors access patterns to identify indicators of compromise, such as mass deletions or rapid file encryption, shifting security from passive retention to active threat mitigation.
- **Workload optimization:** The solution features a sophisticated dynamic data placement engine that optimizes performance and cost within a single namespace. Administrators can configure granular policies to route data to specific storage classes based on object size or tags. An example of this would be automatically directing small, metadata-heavy objects to replicated Non-Volatile Memory Express (NVMe) pools while sending large datasets to erasure-coded high-density Hard Disk Drive (HDD) pools.
- **Versioning:** QuantaStor supports robust S3 bucket versioning integrated with comprehensive lifecycle management (LC) policies. This capability allows for the automated transition of noncurrent versions to lower-cost tiers or their expiration after set intervals, and it includes MFA Delete protection concepts to ensure data recoverability against accidental or malicious user actions.

OPPORTUNITIES

OSNexus has room for improvement in a few decision criteria, including:

- **Auditing:** While the platform maintains comprehensive, immutable audit logs for management operations and supports S3 access logging, it lacks native built-in visualization tools. To achieve deep analysis or real-time threat correlation, users must export logs to external SIEM systems like Splunk or Elasticsearch because the native interface does not offer advanced click-to-view analytics.
- **Storage optimization:** QuantaStor provides effective inline compression and efficient erasure coding, but it does not support global deduplication for object storage pools. The vendor cites performance concerns with upstream Ceph deduplication, but this omission limits storage efficiency for specific backup and archival workloads when deduplication is a primary cost-saving driver.
- **Kubernetes support:** The solution supports Kubernetes via the standard upstream Ceph-CSI driver, enabling essential dynamic

provisioning and snapshots. However, it lacks a proprietary, vendor-specific Kubernetes Operator with deep application-aware extensions or backup logic, lagging behind competitors that offer a more vertically integrated container-native storage experience.

PURCHASE CONSIDERATIONS

OSNexus employs a highly transparent licensing model that stands out for its simplicity, basing costs on raw capacity and subscription duration without the hidden ingress/egress fees often associated with cloud-based alternatives. This straightforward approach, combined with an all-inclusive Enterprise Edition that avoids complex feature-based SKU proliferation, makes budget forecasting notably easier for decision-makers. The solution is well suited for both midsize organizations leveraging the Community Edition and large enterprises requiring petabyte-scale private clouds.

Deployment is flexible, with the solution available as bare metal ISOs, virtual appliances, or cloud images, functioning as a complete unified storage operating system rather than a modular add-on. While the underlying Ceph technology is notoriously complex, QuantaStor's Storage Grid management plane significantly reduces deployment and operational difficulty compared to raw open source alternatives, effectively abstracting the command-line heavy maintenance into a GUI-driven workflow. Furthermore, the platform's extensive hardware compatibility list allows for the repurposing of existing commodity servers, easing migration paths and reducing hardware vendor lock-in.

USE CASES

OSNexus strongly supports large enterprises in regulated sectors like finance, healthcare, and government, leveraging its FIPS 140-2 certification and S3 Object Lock for compliance-heavy workloads. The platform is ideal for implementing massive-scale private clouds, secure backup repositories (validated as Veeam Ready), and active archives. Its Cloud Container and NAS gateway features also make it a strong fit for hybrid cloud architectures requiring seamless data tiering to public cloud targets.

Quantum: Quantum ActiveScale

SOLUTION OVERVIEW

Quantum has evolved from its roots in tape and video surveillance into a comprehensive software-defined data management provider. Its ActiveScale solution is not merely a storage target but a platform designed to manage unstructured data across its entire lifecycle, serving as a forever store that underpins high-performance layers. The solution integrates

distinct media types (NVMe, HDD, and tape) into a single, strongly consistent S3 namespace, collapsing the traditional storage hierarchy. Product SKUs include the Z200 (all-flash), P200 and X200 (hybrid/capacity), and the unique Cold Storage Class architecture utilizing Quantum Scalar tape libraries.

ActiveScale differentiates itself through a focus on economic durability rather than ephemeral feature velocity. It is engineered for risk-averse enterprise architects who prioritize data integrity and extreme durability (19 nines) and predictable economics over bleeding-edge cloud-native integrations. The solution ensures consistency throughout the contract lifecycle, valuing stability and proven reliability for exabyte-scale data lakes.

Quantum is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Quantum scored well on a number of decision criteria, including:

- **Workload optimization:** ActiveScale employs small object aggregation to bundle small files into larger containers, effectively mitigating the metadata overhead and IOPS amplification often associated with small-file archives. Additionally, policy-driven tiering automatically transitions data between high-performance NVMe, HDD, and tape tiers based on administrator-defined S3 Object Lifecycle Management (OLM) policies, optimizing the infrastructure for both performance and cost.
- **Versioning:** The platform supports full S3 bucket versioning, preserving every iteration of an object to safeguard against accidental overwrites or deletions. This capability is tightly integrated with OLM, allowing administrators to define policies that automatically expire older versions or transition them to cheaper storage classes like tape after a set duration.
- **Ransomware protection:** Quantum achieves a distinct advantage by combining logical S3 Object Lock (compliance and governance modes) with the physical security of tape. Data written to the Cold Storage Class can be stored on air-gapped tape cartridges that are physically disconnected from the network, providing an absolute barrier against network-borne cyberthreats that disk-only systems cannot replicate.

OPPORTUNITIES

Quantum has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While ActiveScale supports Kubernetes environments, it relies primarily on generic CSI drivers (S3/NFS)

rather than a proprietary, feature-rich Kubernetes Operator. The solution currently lacks the advanced application-aware features (such as operator-driven lifecycle management or consistent snapshots of complex stateful sets) found in more specialized container-native storage solutions.

- **Auditing:** ActiveScale generates granular audit logs for management and data access operations. However, the platform relies on streaming these logs to external SIEM platforms like Splunk for analysis rather than offering deep, native forensic visualization or behavioral analysis tools within the storage interface itself.
- **Reporting and analytics:** The cloud-based analytics (CBA) dashboard effectively aggregates telemetry for health monitoring and predictive capacity planning. However, the system currently falls short of AIOps-driven autonomous operations. It requires administrative intervention for issue resolution rather than providing AI-based predictive remediation of complex failure scenarios.

PURCHASE CONSIDERATIONS

Quantum offers a transparent licensing model available as a capacity subscription or appliance purchase, avoiding the complexity of unpredictable cloud API and egress fees. The integration of tape as a native storage class enables an extremely low cost for cold data ($0.25-$1.00/TB/month), providing superior TCO visibility for long-term retention. This economic profile makes the solution ideal for large enterprises, CSPs, and MSPs managing petabyte-scale environments.

Deployment flexibility is a hallmark of the offering, with options for turnkey appliances (Z200/P200/X200), software-only licensing for third-party hardware, or fully managed Storage-as-a-Service via Quantum GO. Operational complexity is minimized through a single-pane-of-glass interface that manages both active disk tiers and cold tape tiers uniformly, abstracting the physical library from the user. Furthermore, ActiveScale facilitates cloud repatriation strategies by offering a credible on-prem alternative to AWS Glacier, with the unique ability to replicate data directly to or from public cloud cold tiers.

USE CASES

ActiveScale supports industry verticals requiring massive, durable data retention, including media and entertainment (M&E), life sciences, and the federal government sector. In M&E, it serves as a content repository where raw footage remains accessible via S3 while stored economically on tape. As a Platform Play, it supports broad use cases ranging from active archives to regulatory compliance vaults, effectively functioning as a private cloud in a box that mirrors the tiered storage services of public cloud providers.

Qumulo: Qumulo Core

SOLUTION OVERVIEW

Qumulo focuses on enterprise-scale unstructured data management, rejecting the traditional boundaries between on-prem and cloud infrastructure. Its Run Anywhere philosophy allows the same software stack to run on proprietary appliances, certified commodity hardware from vendors like HPE and Cisco, and public cloud infrastructure. Recently, the vendor has aggressively pivoted toward cloud-native architectures, refactoring its platform to decouple compute from storage in its Cloud Native Qumulo (CNQ) and Azure Native Qumulo (ANQ) offerings.

The primary solution, Qumulo Core, is a software-defined distributed file system that natively supports S3, NFS, SMB, and FTP protocols. Rather than acting as a simple gateway, Qumulo Core writes data to a unified block store, allowing files to be accessed as objects and vice versa without data duplication. This portfolio includes the Qumulo Nexus control plane for global management and Run Anywhere for software-only deployments.

Qumulo is positioned in the Innovation hemisphere, reflecting its rapid architectural shifts and frequent release cadence. The solution will look and feel different over the contract lifecycle. Qumulo delivers an aggressive roadmap and is dynamically responsive to user needs, valuing rapid advancement and frequent releases. It puts primary emphasis on emerging features and rapid development to address gaps, such as the move to disaggregated cloud architectures.

Qumulo is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Qumulo scored well on a number of decision criteria, including:

- **Public cloud integration:** Qumulo delivers a seamless hybrid cloud experience through Cloud Native Qumulo (CNQ) and Azure Native Qumulo (ANQ). These services do not merely run virtual machines in the cloud; they leverage the cloud's native object storage (AWS S3, Azure Blob, Google Cloud Storage, and Oracle Object Storage) as the backend persistence layer while using ephemeral compute for file system logic. This architecture decouples compute from storage, allowing users to scale performance independently of capacity. Furthermore, the global namespace unifies these instances with on-prem clusters, and the shift feature allows bidirectional movement of data between Qumulo's proprietary format and open S3 objects.

- **Reporting and analytics:** The platform provides exceptional real-time visibility through the Qumulo Nexus dashboard and local UI. Unlike competitors that rely on delayed reporting, Qumulo allows administrators to see exactly which client IP is consuming the most bandwidth, which directory is growing fastest, and what latency distribution looks like instantaneously. The system aggregates telemetry to identify hardware failure trends and capacity exhaustion dates months in advance, drilling down from a global view to specific file throughput.
- **Ransomware protection:** Qumulo supports S3 Object Lock in compliance mode, ensuring that a retention period is set during which an object cannot be overwritten or deleted by any user, including root. This native immutability is complemented by snapshot locking for file system protection. Additionally, through integrations with partners like Varonis, the system facilitates active monitoring of access patterns to detect attack signatures and trigger automated lockdown responses.

OPPORTUNITIES

Qumulo has room for improvement in a few decision criteria, including:

- **Workload optimization:** While Qumulo effectively uses its NeuralCache technology to analyze access patterns and prefetch data into RAM or NVMe, it lacks granular administrator controls. The system operates as a massive distributed cache, but explicitly configurable QoS policies (such as manually pinning a specific S3 bucket to NVMe flash or setting hard IOPS limits for a specific tenant) are less capable than competing block or object storage systems.
- **Versioning:** The platform supports standard S3 bucket versioning, allowing users to retain, list, and restore older versions of objects when new writes occur. However, the implementation misses the nuanced lifecycle management features found in AWS S3. Complex transition rules (such as automatically moving specific versions to deep archive tiers after a set period) are limited because Qumulo often acts as the single tier of storage unless using the external shift feature.
- **Storage optimization:** Qumulo employs a highly adaptive erasure coding scheme that achieves high usable capacity efficiencies (up to 85%), avoiding the overhead of rigid RAID. In the cloud, its CNQ and ANQ offerings further optimize storage through built-in compression and the consolidation of files into fewer objects to

minimize storage and transaction costs. However, the solution still lacks a global, inline deduplication consistent across deployment models.

PURCHASE CONSIDERATIONS

Licensing transparency is a key differentiator for Qumulo. The Qumulo One program offers a flat, predictable subscription model that is portable between on-prem and cloud environments. If a customer purchases a license for on-prem capacity and later migrates data to AWS, the license transfers without penalty, eliminating the double-payment issue common in hybrid migrations. Pricing for cloud-native offerings is similarly transparent, separating performance costs from capacity costs.

The solution is targeted primarily at the Global 2000 and large enterprise sector rather than SMBs. Minimum node requirements for erasure coding efficiency and the cost structure of enterprise-grade hardware typically place it outside the budget of small businesses. It is well suited for organizations requiring petabyte-scale data sets for whom management overhead reduction is critical.

Deployment is highly flexible, supporting physical appliances from HPE, Cisco, and Supermicro, as well as software-only deployments on qualified commodity hardware. Management is simplified via Qumulo Nexus, a SaaS-based control plane that unifies disparate clusters, making the solution easier to manage than many legacy parallel file systems.

USE CASES

Qumulo supports industries with massive unstructured data requirements, including media and entertainment for 4K/8K rendering, life sciences for genomics, and healthcare for PACS imaging. Specific use cases include high-performance analytics and active archives for which data must remain immediately accessible. Qumulo correlates these use cases into a holistic data management environment, unifying the lifecycle of data from creation via high-performance file protocols to processing and distribution via S3, without requiring distinct storage silos.

Quobyte

SOLUTION OVERVIEW

Quobyte is an SDS vendor founded by former Google infrastructure engineers, focusing on high-performance object storage and data center file systems (DCFS). The company addresses the convergence of HPC and enterprise IT, specifically targeting fast and vast

data workloads such as AI training and genomic sequencing. Quobyte differentiates itself by offering a unified data plane that provides simultaneous, consistent access to data via file (POSIX, NFS, SMB) and object (S3) protocols without the need for gateways or data migration.

Quobyte is a single product that runs on commodity x86 or ARM servers, allowing organizations to build exabyte-scale clusters with a mix of media types. It is not part of a modular suite but rather a comprehensive software stack that handles metadata, data services, and management capabilities within a shared-nothing architecture. Quobyte's strategy prioritizes architectural purity and performance over legacy compatibility, eschewing standard appliances for a hyperscaler-inspired software model.

Quobyte is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Quobyte scored well on a number of decision criteria, including:

- **Workload optimization**: Quobyte excels in AI/ML environments, verified by MLPerf Storage benchmarks against which it demonstrated high GPU utilization (over 90%) for rigorous image segmentation tasks. The architecture supports remote direct memory access (RDMA) and parallel I/O, allowing clients to stream data from all available storage nodes simultaneously. This design bypasses the CPU for data transfers, significantly reducing latency for throughput-intensive workloads like deep learning and high-frequency trading.
- **Versioning:** The platform implements a unified versioning system that bridges the gap between file and object storage. Unlike typical object stores from which versions are accessible only via APIs, Quobyte exposes S3 object versions as file snapshots accessible through a hidden directory in the file system. This capability simplifies recovery workflows, enabling users to browse and restore previous versions of an S3 object using standard file explorers or terminal commands.
- **Ransomware protection:** Quobyte provides robust defense mechanisms through immutable storage policies. It supports S3 Object Lock in both compliance and governance modes, which maps directly to file system WORM policies. Administrators can configure immutable snapshots that remain read-only even if the live file system is compromised, ensuring a clean recovery point is always available.

OPPORTUNITIES

Quobyte has room for improvement in a few decision criteria, including:

- **Auditing:** While Quobyte logs every operation, including reads, writes, and metadata changes, it lacks a native built-in visualizer for these audit trails. The solution streams events to external targets like Kafka for integration with SIEM tools, but organizations must rely on third-party ecosystems (like Splunk or ELK) to visualize, analyze, or alert on this data, adding a layer of operational complexity.
- **Public cloud integration:** Quobyte runs natively on public cloud infrastructure (AWS, GCP, Azure) and supports hybrid tiering, but it operates as a self-managed software layer on top of IaaS rather than a fully managed SaaS offering. While the Data Mover service manages hydration between on-prem and cloud S3 targets, the lack of a zero-touch, vendor-managed cloud service requires users to maintain the underlying cloud compute and storage resources themselves.
- **Reporting and analytics:** The system eschews proprietary reporting tools in favor of open observability standards, exposing metrics via a native Prometheus exporter. Although Quobyte provides prebuilt Grafana dashboards, it does not offer a self-contained, native user behavior analytics (UBA) dashboard. Consequently, administrators must maintain the external visualization stack to perform long-term trend analysis and capacity planning.

PURCHASE CONSIDERATIONS

Quobyte offers a transparent, capacity-based licensing model that avoids the complexity of per-user fees or nickel-and-dime charges for API requests and data egress. This predictability is a significant advantage for AI/ML workloads involving massive data rereading (epochs), for which public cloud costs can spiral. The software-only delivery model supports “deployment anywhere,” allowing enterprises to utilize existing commodity hardware, including the latest x86 or ARM processors, without waiting for appliance validation.

The solution is best suited for large enterprises and CSPs with significant technical in-house capability, particularly those managing petabyte-scale data lakes. Deployment requires a greenfield approach or a planned migration, as the system utilizes a unique DCFS architecture rather than overlaying existing storage. While the API-first design enables high automation (smartphone simplicity), effective management at scale still necessitates Linux and networking expertise, making it less ideal for SMBs looking for a black box appliance.

USE CASES

Quobyte primarily supports performance-intensive verticals including AI, life sciences, and M&E. In AI training pipelines, it saturates GPU clusters via RDMA, handling the massive bandwidth required for LLMs. For M&E, it manages large sequential I/O streams essential for 4K/8K editing and rendering. It unifies these workloads, serving as a single backend for block (via Cinder/CSI), file (NFS/POSIX), and object (S3) storage, eliminating data silos.

Scality: RING

SOLUTION OVERVIEW

Scality is an SDS vendor focused on enabling large-scale enterprises and service providers to manage unstructured data without hardware lock-in. The company operates independently, delivering its flagship solution, Scality RING, as a platform designed to run on standard commodity x64 servers. RING is a unified solution that combines native file (NFS/SMB) and object (S3) interfaces within a single distributed system, allowing organizations to consolidate disparate data silos. The architecture supports MultiScale growth, meaning capacity and performance can be scaled independently to suit varying workload requirements.

Scality is positioned in the Innovation hemisphere, reflecting its aggressive development roadmap and pivot toward high-performance workloads. The solution will look and feel different over the contract lifecycle as the vendor delivers rapid advancements to address modern needs. Scality actively pushes beyond traditional archival use cases, emphasizing performance-oriented features like the RING XP configuration for AI and open source standards for Kubernetes. Primary emphasis is on enabling new capabilities for AI data lakes and cloud-native integration while maintaining the stability required for exabyte-scale deployments.

Scality is positioned as a Leader and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Scality scored well on a number of decision criteria, including:

- **Workload optimization:** Scality addresses diverse performance needs through its MultiScale architecture, which allows compute (connectors) and storage resources to scale independently. The vendor recently introduced RING XP, an all-NVMe flash configuration that bypasses standard S3 overhead to deliver microsecond-level latencies (500µs read/600µs write) specifically for AI model training and small-file workloads.

- **Storage optimization:** The solution maximizes efficiency through its Storage Accelerator, a tiering mechanism that lands hot data on flash before migrating it to cost-effective HDDs, compressing data during the transition. The platform utilizes advanced erasure coding schemas, like geo-distributed EC, to ensure high durability with significantly lower storage overhead compared to traditional replication.
- **Ransomware protection:** Scality implements a comprehensive CORE5 defense strategy covering API, data, storage, architecture, and geographic levels. Key capabilities include immutable S3 Object Locking (compliance and governance modes), data-at-rest encryption with external Key Management Service (KMS) integration, and an architecture that never overwrites data in place, inherently protecting against encryption attacks.

OPPORTUNITIES

Scality has room for improvement in a few decision criteria, including:

- **Versioning:** The platform supports standard Amazon S3 versioning APIs, allowing for the preservation and retrieval of previous object states. While functional and fully compliant with industry standards for recovering from accidental deletions or overwrites, the feature set represents a baseline implementation rather than a unique differentiator with advanced management capabilities.
- **Public cloud integration:** Scality excels at hybrid cloud tiering and bidirectional replication, but it does not offer a cloud-native image (AMI) for running the RING software itself directly on public cloud infrastructure (AWS, Azure) as a primary deployment model. Organizations seeking to lift and shift the same storage software stack into the public cloud for processing may find this limiting compared to peers with native cloud marketplace instances.
- **Reporting and analytics:** The solution provides deep visibility through its cloud monitor console and Maestro federated management plane. However, the platform currently relies on external tools or manual analysis for some predictive insights, as the fully autonomous, embedded AI agents (Guardian) for proactive issue prevention and self-remediation are currently roadmap items slated for 2026.

PURCHASE CONSIDERATIONS

Scality employs a transparent, capacity-based subscription model that is inclusive of all features, simplifying licensing for buyers. Pricing is based on usable capacity rather than raw storage or number of nodes, which encourages efficient data protection strategies without financial penalty. This approach avoids the hardware tax often

associated with appliance-based vendors, as users can source competitive hardware bids for the underlying infrastructure.

The solution is software-only and designed for deployment on bare metal Linux servers. This requires a degree of architectural planning that may be more complex than plugging in a turnkey appliance, though it offers greater long-term flexibility. The buyer profile leans heavily toward CSPs, MSPs, and large enterprises (Global 2000) that require massive scale and specific customizability. While zero-touch operations are a goal, managing a large-scale distributed system like RING typically requires a mature IT operations team.

USE CASES

Scality primarily supports large enterprises and CSPs in verticals such as banking, healthcare, and media. Specific use cases include AI data lakes, where RING XP provides the necessary throughput for GPU saturation; sovereign cloud infrastructure for regional service providers; and massive-scale compliance archives for sectors with strict regulatory retention requirements.

Seagate: Lyve Cloud Object Storage

SOLUTION OVERVIEW

Seagate, historically dominant in the HDD market, has expanded into the systems and services layer to capture value across the data lifecycle. The vendor's object storage strategy leverages vertical integration, combining its own media manufacturing with cloud services to offer a disruptive cost structure. Recent efforts focus on the mass data problem, positioning Seagate as a neutral repository for petabyte-scale workloads.

Lyve Cloud object storage is a storage-as-a-service (STaaS) offering designed for multicloud accessibility and predictable economics. It operates as a standalone S3-compatible object store but is often deployed alongside Seagate's Exos systems for on-prem needs. The solution targets large enterprises requiring mass capacity for backup, archive, and data lakes.

The solution will look and feel largely the same throughout the contract lifecycle, prioritizing operational predictability and a simplified economic model. Seagate values steady platform reliability, consistent service delivery, and a vertically integrated cost structure over rapid feature expansion, an approach suited to organizations seeking stable, long-term capacity commitments.

Seagate is positioned as a Challenger and Fast Mover alone in the Maturity/Feature Play quadrant of the object storage Radar chart.

STRENGTHS

Seagate scored well on a number of decision criteria, including:

- **Public cloud integration**: Lyve Cloud is architected to eliminate egress fees and API charges, differentiating it significantly from hyperscalers. This design enables a true multicloud workflow whereby data stored in Lyve Cloud can be accessed by compute resources in AWS, Azure, or GCP without financial penalty, fostering genuine cloud neutrality.
- **Auditing:** The platform provides robust logging capabilities, generating detailed JSON logs for all Console and S3 API activities. These logs facilitate deep security visibility and compliance monitoring, with documented integration for enterprise SIEM tools like Splunk and AWS CloudWatch.
- **Ransomware protection:** Seagate supports S3 Object Lock in both governance and compliance modes, effectively logically air-gapping data to prevent deletion or modification. This immutability is critical for regulatory compliance and protecting backup targets against malicious encryption attacks.

OPPORTUNITIES

Seagate has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While basic S3 connectivity is available, the solution lacks a general availability (GA) COSI driver. Integrations currently rely on standard S3 connectors or FUSE adapters rather than native Kubernetes storage orchestration.
- **Storage optimization:** The architecture prioritizes mass capacity economics but lacks native inline global deduplication within the object store itself. Data reduction relies on client-side software or backup partners, which limits storage efficiency for workloads that are not preoptimized before ingestion.
- **Reporting and analytics:** The native console provides basic capacity reporting, but advanced insights are offloaded to the separate Lyve Cloud Analytics Platform. While this external platform is capable, the core object storage service lacks deeply integrated, granular native analytics, requiring customers to adopt an additional DataOps layer for detailed visibility.

PURCHASE CONSIDERATIONS

Seagate's pricing strategy is a primary purchase driver, characterized by exceptional transparency. The flat-rate capacity model with zero egress or API fees eliminates the sticker shock that's common with hyperscale providers, making TCO highly predictable.

The solution is well suited for large enterprises managing petabyte-scale datasets, particularly in media, healthcare, and backup use cases. While Seagate can support SMB workloads via its Lyve Management Portal, the solution is primarily optimized for mass capacity; therefore, organizations seeking a small, general-purpose bucket may find its architecture more robust than their needs require.

Deployment is straightforward as a SaaS offering, removing infrastructure management overhead. However, for on-prem needs, it requires the adoption of Exos systems. Migration is facilitated by standard S3 compatibility, and migration is further simplified by a native Global Endpoint, which allows customers to manage and access data across all global regions through a single namespace. The ecosystem is heavily weighted toward backup and archive partners like Veeam and Commvault, ensuring strong support for data protection workflows.

USE CASES

Seagate supports verticals with massive data generation needs, such as M&E for active archives and healthcare for genomics data repositories. The solution excels in use cases requiring high-durability storage with frequent access from multiple clouds, such as data lakes and backup repositories. Seagate focuses specifically on solving the cost and gravity friction of mass storage rather than attempting to be a general-purpose compute cloud.

SoftIron: Hypercloud

SOLUTION OVERVIEW

SoftIron designs and manufactures task-specific storage and compute appliances, distinguishing itself with a true private cloud philosophy that reintegrates hardware and software to eliminate complexity. By controlling the entire stack from motherboard design to the source-compiled operating system, the vendor offers a secure, turnkey infrastructure solution aimed at organizations seeking to repatriate workloads from public clouds without sacrificing operational elasticity. Recent strategic moves include the introduction of industry-first post-quantum cryptography (PQC) capabilities, reinforcing a focus on data sovereignty and long-term security for regulated sectors.

HyperCloud is a single, unified platform rather than a loose collection of products. It converges block, file, and object storage with compute virtualization and networking into a self-healing fleet. The HyperCloud ecosystem is composed of specialized Interconnect, Storage, and Compute (including GPU) nodes, which are deployed as physical appliances to form a unified, self-healing fabric.

The solution will look and feel largely the same throughout the contract lifecycle, prioritizing stability and continuity. SoftIron values incremental improvement, consistent user experience, and assured compatibility over breakneck advancement.

SoftIron is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

SoftIron scored well on a number of decision criteria, including:

- **Auditing:** HyperCloud delivers superior visibility through granular logging of all administrative and data access events. The system captures detailed audit trails (including user identity, object access, and outcomes) and supports real-time export to external SIEM platforms like Splunk and QRadar via standard syslog protocols for compliance and forensics.
- **Versioning:** The platform supports standard S3 bucket versioning, enabling the preservation and retrieval of multiple object variants. This capability allows users to recover from accidental deletions or overwrites and includes lifecycle rules to manage the expiration of older versions, providing a critical safety net for data durability.
- **Ransomware protection:** SoftIron addresses data security through S3 Object Lock, supporting both compliance and governance modes. This WORM capability ensures that backup data and critical archives remain immutable and cannot be encrypted or deleted by attackers, including those with root privileges.

OPPORTUNITIES

SoftIron has room for improvement in a few decision criteria, including:

- **Public cloud integration:** Designed as a replacement for public cloud infrastructure, HyperCloud lacks native bidirectional tiering to hyperscalers like AWS S3 or Azure Blob. The solution functions primarily as an on-prem silo, limiting its utility for organizations requiring seamless hybrid data mobility or cloud bursting for object storage specifically.
- **Storage optimization:** While the platform employs efficient erasure coding and copy-on-write mechanisms, it lacks global inline deduplication and compression for object data. This limitation may result in lower effective capacity utilization compared to competitors that offer advanced data reduction pipelines for repetitive workloads.
- **Reporting and analytics:** The Glasshouse management dashboard provides descriptive metrics on cluster health and capacity but

lacks advanced predictive analytics. The system does not currently feature AI-driven forecasting or anomaly detection to predict capacity crunches or performance bottlenecks before they occur.

PURCHASE CONSIDERATIONS

SoftIron distinguishes itself with a transparent licensing model that eliminates the unpredictable costs often associated with cloud storage. The solution is sold as capital equipment (CapEx) or via an OpEx subscription, with no data egress fees or per-request API charges, allowing for precise TCO modeling.

The primary buyer profile includes large enterprises and MSPs looking to consolidate workloads or build sovereign clouds. Deployment is strictly appliance-based; there is no software-only or virtual appliance option, as the hardware is integral to the security and performance guarantees. Professional services are often leveraged for initial fleet setup, though the fleet intelligence automation simplifies ongoing node additions. Migration is typically a repatriation effort, moving data from public clouds or legacy SANs onto the HyperCloud fabric.

USE CASES

SoftIron supports verticals requiring strict data sovereignty and security, such as government, defense, and finance, leveraging its secure provenance manufacturing and PQC features. Primary use cases include cloud repatriation for consolidating massive data sets and modernizing legacy virtualization stacks onto a single fabric. The platform supports these diverse workloads by unifying S3 object storage with block and file services, allowing it to serve as a general-purpose private cloud rather than just a dedicated storage array.

Spectra Logic: BlackPearl Object Gateway, S3 Hybrid object storage

SOLUTION OVERVIEW

Spectra Logic has evolved from its roots in robotic tape libraries to become a specialized provider of data management and object storage solutions, focusing on the hybrid cloud intersection of active archiving and long-term data preservation. The core solution, BlackPearl Object Gateway software, operates as a converged storage architecture that integrates flash, disk, and tape behind a unified S3-compatible interface. This portfolio approach allows the vendor to address the complete data lifecycle, managing data from hot ingest on flash to cold air-gapped tape without requiring the application to understand the underlying media complexities.

Spectra Logic prioritizes a hybrid-first strategy, fundamentally challenging the all-cloud orthodoxy by enabling organizations to maintain authoritative Golden Copies on-prem while leveraging public cloud compute elasticity. The vendor is positioned in the Innovation hemisphere because its roadmap emphasizes dynamic software advancements like On-Premises Glacier API translation, which aggressively redefines how legacy media interacts with modern cloud-native workflows.

Spectra Logic is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Spectra Logic scored well on a number of decision criteria, including:

- **Public cloud integration:** The solution creates a federated global namespace across on-prem sites and public cloud regions. Uniquely, the platform can intercept standard Amazon S3 Glacier API commands and translate them into robotic tape operations, allowing organizations to repatriate cloud archive workflows to local infrastructure without code changes. The system supports bidirectional synchronization and cloud bursting, enabling data to be moved to the cloud for transient compute tasks and then deleted to avoid long-term storage fees.
- **Versioning:** Spectra Logic's implementation of S3 compatibility includes robust support for object versioning, which serves as the foundation for its data immutability features. By leveraging S3 Object Lock in both governance and compliance modes, the platform ensures that specific object versions cannot be deleted or overwritten until retention periods expire, meeting strict regulatory requirements like SEC 17a-4. This capability is integrated with the vendor's lifecycle policies, ensuring versions are correctly tiered to appropriate media for cost-effective retention.
- **Ransomware protection:** The vendor offers an attack-hardened defense-in-depth architecture that combines logical immutability with a true physical air gap. Unlike disk-only competitors that rely on network segmentation, Spectra Logic's integration with their tape libraries physically decouples media from the network, rendering it inaccessible to remote attackers. This is complemented by trigger-based snapshots that preserve data states upon detecting anomalous activity and strict administrative lockdowns that require local console access for critical changes.

OPPORTUNITIES

Spectra Logic has room for improvement in a few decision criteria, including:

- **Kubernetes support:** Spectra Logic's Kubernetes support is intentionally object-centric, focusing on native S3 compatibility for containerized workloads rather than block-level persistent volume provisioning. While this aligns with the product's design as a specialized object store, the platform lacks a native CSI driver. This limits its utility for teams requiring automated volume lifecycle management or the plug-and-play persistent storage integration found in solutions from MinIO or Dell.
- **Storage optimization:** While the system is highly efficient regarding power and cooling, it lacks the inline global deduplication capabilities often found in dedicated backup appliances. For datasets with high redundancy, such as virtual machine backups, this omission can result in higher capacity consumption compared to deduplication-centric peers. The vendor prioritizes low-cost media (tape) over compute-intensive deduplication algorithms to achieve TCO targets.
- **Reporting and analytics:** Native metadata management is functional, supporting standard S3 tags for lifecycle policies, but lacks deep content indexing or advanced analytics capabilities out of the box. Organizations requiring complex global search or the ability to query metadata attributes (such as extracting text or genomic headers) must rely on the generally available Mediaflux integration with partner Arcitecta. This reliance on a third-party software stack prevents the core platform from offering a unified, native database-like query experience.

PURCHASE CONSIDERATIONS

Spectra Logic's licensing model for the BlackPearl software is transparent, based on managed capacity rather than complex API transaction fees, which contrasts favorably with the unpredictable egress charges of public cloud providers. The hardware deployment typically involves the BlackPearl appliance acting as a gateway to high-density expansion disk chassis and/or robotic tape libraries, presenting a converged platform rather than a loose collection of nodes.

This solution is best suited for medium to large enterprises, particularly in M&E, government, and research sectors because their sheer volumes of data make public cloud storage economically prohibitive. Deployment can be greenfield or integrated into existing S3 workflows, but the unique tape-as-object architecture means migration often involves strategic decisions about data tiering policies rather than a simple lift-and-shift. While the user interface abstracts the complexity of robotics, organizations should consider the total resource consumption (including power, cooling, and floorspace) of tape libraries compared to dense flash arrays.

USE CASES

Spectra Logic does well in M&E workflows, where high sequential throughput supports massive video ingest and active archiving without dropped frames. It is also heavily utilized in HPC and genomic research, serving as the repository for exabyte-scale datasets that require long-term preservation at the lowest possible cost. While highly specialized for these data-intensive fields, the platform is architected for broad enterprise preservation, acting as a central data mover that orchestrates content between expensive primary flash and economical deep storage.

Storj: Storj Object Storage

SOLUTION OVERVIEW

Storj has evolved from a decentralized storage network into a broader cloud infrastructure provider, challenging hyperscalers with its decentralized physical infrastructure networks (DePIN) architecture. Rather than relying on centralized data centers, Storj aggregates capacity from a global network of independent nodes, utilizing erasure coding to ensure durability and availability without the distance tax or complex replication schemes of traditional providers. In the last year, the vendor has aggressively expanded its portfolio through the acquisitions of Valdi (high-performance GPU compute) and PetaGene (the technology behind cunoFS), pivoting from a storage-only focus to a comprehensive compute-and-storage platform.

Storj object storage integrates S3-compatible object storage, high-performance file access via Object Mount, and emerging compute capabilities. The vendor is positioned in the Innovation hemisphere, reflecting its disruptive architectural approach and rapid roadmap development (specifically its move to unify storage with compute-over-data workflows for AI and media use cases). Storj prioritizes breaking performance and cost barriers over maintaining legacy code bases, using a distributed model that inherently offers multiregion durability.

Storj is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Storj scored well on a number of decision criteria, including:

- **Kubernetes support:** The platform utilizes the cunoFS CSI Driver to enable dynamic provisioning of persistent volumes (PVs) backed directly by Storj buckets. This driver uses a DaemonSet architecture to mount buckets as local file systems within pods, supporting ReadWriteMany access modes and enabling data portability for stateful workloads like AI training.

- **Versioning:** Storj provides robust S3-compatible object versioning, allowing users to maintain multiple variants of an object in the same bucket and restore previous states using standard API calls. This capability is tightly integrated with the platform's immutable storage features, facilitating the protection of historical versions against accidental deletion or overwriting.
- **Ransomware protection:** The solution supports S3 Object Lock in compliance mode, ensuring data cannot be deleted or overwritten by any user, including the root account, until a specified retention period expires. This immutability, combined with the distributed nature of the network, provides a resilient defense against encryption attacks and malicious deletion.

OPPORTUNITIES

Storj has room for improvement in a few decision criteria, including:

- **Workload optimization:** While the platform delivers high parallel throughput, it lacks native policy-driven automated tiering to move data between hot and cold media based on access frequency. Users cannot configure QoS controls to prioritize specific tenants or guarantee IOPS for critical workloads.
- **Storage optimization:** Storj relies entirely on erasure coding for storage efficiency and does not offer inline compression or global deduplication. While erasure coding is efficient for durability, the lack of native data reduction features limits the platform's effectiveness for workloads with high data redundancy, such as backup targets that do not perform client-side deduplication.
- **Reporting and analytics:** The console provides basic visibility into stored bytes, egress usage, and object counts, but it lacks built-in predictive analytics for capacity forecasting or performance trend analysis. Users seeking deep insights or anomaly detection must currently export raw logs to external SIEM or observability platforms.

PURCHASE CONSIDERATIONS

Storj distinguishes itself through a transparent pricing model that avoids the complex API charges and multilayered fee structures common among hyperscalers. While the platform now offers a small set of workload-oriented storage tiers based on geography and egress requirements, it avoids the performance-based or access-frequency policies typical of the market. The use of flat-rate pricing for storage and egress options continues to simplify financial forecasting, making it particularly attractive for SMBs and enterprises with unpredictable data retrieval patterns. Unlike modular solutions that require piecemeal licensing, the Storj offering is consumed

primarily as a managed service, reducing deployment complexity for the core storage layer.

Deployment does not require proprietary hardware. The network runs on commodity hardware provided by node operators, while customers interact via standard S3 APIs or the Gateway-MT. Migration is facilitated by standard tools like Rclone and Veeam, though the unique decentralized architecture means it functions best as a greenfield or tiering target rather than a direct drop-in replacement for low-latency block storage. The addition of Object Mount technology significantly eases adoption for organizations requiring POSIX file system semantics alongside object storage.

USE CASES

Storj supports M&E workflows by enabling global, high-performance file access for editing and rendering without the need for regional replication. For AI/ML and data-intensive workloads, the platform integrates distributed CPU and GPU compute resources, allowing models to access datasets directly from the storage layer and minimizing data movement costs. Storj also serves as a cost-effective, immutable target for backup and archive data, leveraging its inherent multiregion durability to simplify disaster recovery strategies.

StorONE: S1 Enterprise Storage Platform

SOLUTION OVERVIEW

StorONE is an SDS vendor focused on maximizing hardware efficiency through a complete rewrite of the traditional storage IO stack. Its flagship offering, the S1 Enterprise Storage Platform, is a unified solution that consolidates block, file, and object workloads onto a single layer of commodity hardware, eliminating the need for fragmented storage silos. Unlike competitors that bolt object storage onto legacy architectures via gateways, StorONE utilizes a virtual storage container abstraction that allows S3 workloads to run natively alongside high-performance block volumes on the same underlying media pool.

The solution is delivered as a single software platform, licensed via a TRUprice model based on the number of drive slots rather than capacity, ensuring cost transparency and scalability. StorONE's strategy emphasizes a Maturity approach, prioritizing stability, consistent user experience across protocols, and the ability to leverage standard x86 servers and hybrid media (NVMe + HDD) for optimal price and performance. This positioning appeals to enterprises seeking to modernize infrastructure without the lock-in of proprietary appliances.

StorONE is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

StorONE scored well on a number of decision criteria, including:

- **Workload optimization:** The S1 platform features TierONE, an intelligent, real-time auto-tiering engine that optimizes data placement across mixed media types (NVMe, SSD, HDD) within the same volume. Unlike legacy tiering that moves data daily, TierONE operates continuously, ensuring incoming object writes land on flash for immediate performance while transparently migrating cooling data to high-capacity HDDs. This allows the solution to deliver all-flash performance profiles for active data lakes while maintaining the economic benefits of hybrid storage.
- **Storage optimization:** StorONE replaces traditional RAID with vRAID (virtual RAID), a high-performance erasure coding implementation that significantly improves drive utilization and recovery speeds. The system can utilize up to 90% of available drive capacity without the performance degradation typical of legacy filesystems. Furthermore, vRAID enables declustered rebuilds that can restore a failed 18TB drive in under three hours, drastically reducing the vulnerability window compared to standard RAID architectures.
- **Ransomware protection:** The platform employs a recovery-first strategy centered on its SnapONE technology, which supports unlimited immutable snapshots. These snapshots are protected by vSnapLock, creating a logical air gap that prevents deletion even by compromised administrator accounts. Combined with anomaly detection that alerts on suspicious write patterns, this architecture allows organizations to recover petabyte-scale object stores to a pre-attack state in minutes, offering a superior RTO compared to restoring from external backups.

OPPORTUNITIES

StorONE has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While the platform provides a compliant CSI driver for block storage provisioning, its support for object storage in Kubernetes is currently limited to standard S3 API consumption without a native COSI driver. Dynamic provisioning and lifecycle management of object buckets within Kubernetes clusters currently rely on external automation or APIs rather than a fully integrated, Kubernetes-native control plane.
- **Auditing:** The solution offers standard audit logging for system access, configuration changes, and user events, but it lacks advanced, granular content analysis capabilities. While the platform

supports HIPAA-compliant audit workflows and content querying via ONEai, the native auditing engine lacks real-time streaming integration with major SIEM platforms like Splunk or Microsoft Sentinel. Current operations rely on manual log exports or third-party visualization tools for advanced security forensics and real-time anomaly alerting.

- **Reporting and analytics:** Although the S1:LMA (logging, monitoring, alerting) tool provides essential visibility into performance and capacity, the platform lacks the deep, predictive AI-driven analytics found in some market leaders. While capacity prediction is available, the solution relies on integrations with external tools like Grafana for advanced visualization and does not yet offer a comprehensive, native AIOps suite for proactive infrastructure management.

PURCHASE CONSIDERATIONS

StorONE differentiates itself with a transparent TRUprice licensing model that charges based on the number of drive slots populated rather than the total capacity or specific media type used. This approach prevents sticker shock as drive densities increase. A 20TB drive costs the same to license as a 4TB drive, encouraging the use of high-density media.

The solution is highly flexible in deployment, available as software-only for installation on qualified commodity servers or as a bundled solution through partners like HPE and Seagate. This allows buyers to leverage existing hardware relationships and avoid vendor lock-in. The platform is well suited for midsize to large enterprises looking to consolidate backup, archive, and production workloads, though the DIY nature of SDS on commodity hardware may introduce deployment complexity compared to turnkey appliances from legacy vendors.

USE CASES

StorONE supports data-intensive verticals such as healthcare and M&E by delivering a unified high-performance tier for active archives and PACS imaging. The platform excels in AI data lakes, utilizing NVMe-oF and high-speed S3 to minimize wait times for GPU-intensive training and inference workloads. Additionally, it serves as a robust target for high-performance backup solutions like Veeam and Commvault, offering rapid ingest and instant VM recovery capabilities. This versatility enables DevOps environments to consolidate persistent block volumes and S3 buckets onto a single, resource-efficient storage pool, validating its position as a comprehensive Platform Play.

VAST Data: VAST AI OS

SOLUTION OVERVIEW

VAST Data has repositioned itself from a storage vendor to the provider of the VAST AI Operating System, a unified data platform designed explicitly for the AI era. The solution targets high-end enterprises and neocloud service providers, rejecting the traditional shared-nothing architecture in favor of a disaggregated shared-everything (DASE) model that decouples compute from state. This architecture consolidates file, object, and database services into a single repository, eliminating the need for tiering by leveraging NVMe-over-Fabrics and storage class memory (SCM) to deliver all-flash performance at archive economics. The platform is available as a physical appliance (via manufacturing partners), a software-only option for qualified hardware, or as a service on major public clouds (AWS, Azure, GCP).

VAST Data effectively collapses the silos of file systems, object stores, and databases into a unified AI Factory. The vendor is firmly positioned in the Innovation hemisphere, characterized by an aggressive roadmap that includes the integration of vector database capabilities directly into the storage layer via the InsightEngine.

VAST Data is positioned as a Leader and Outperformer in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

VAST Data scored well on a number of decision criteria, including:

- **Workload optimization:** The platform is explicitly engineered for high-performance AI and HPC workloads, supporting NVIDIA GPUDirect Storage (GDS) to bypass CPU bottlenecks and deliver data directly to GPU memory. The DASE architecture eliminates east-west traffic, allowing for massive parallelism and linear performance scaling that rivals high-performance block storage, validated by top-tier benchmarks.
- **Storage optimization:** VAST employs a revolutionary similarity-based data reduction technology that compresses blocks that are merely similar rather than identical, achieving reduction ratios that make all-flash storage economically competitive with HDD-based archives. This global reduction operates across the entire namespace without the performance penalties associated with traditional deduplication.
- **Ransomware protection:** The solution features indestructible snapshots, which provide a robust layer of immutability that withstands even compromised root or admin credentials. These snapshots are protected by time locks and require a multifactor

authentication process involving a support token exchange to delete before expiration, effectively neutralizing internal and external threats.

VAST Data is classified as an Outperformer because of its rapid pace of innovation, particularly the introduction of the InsightEngine and SyncEngine, which transform the system from a passive repository into an active data processing engine. Integrating vector embedding generation directly into the storage path is a leap that positions the vendor to lead the market in simplifying generative AI infrastructure.

OPPORTUNITIES

VAST Data has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While VAST provides a generally available COSI driver, the solution currently faces scaling limitations that complicate edge deployments, as clusters typically do not scale to less than 300 TB of usable capacity. Additionally, the vendor acknowledges that current VAST on Cloud implementations can be complex and lack the elasticity found in cloud-native specific offerings.
- **Versioning:** The platform supports standard S3 bucket versioning to retain and restore object iterations, but it currently lacks support for the S3 MFA Delete API feature. While VAST mitigates this with indestructible snapshots, the absence of strict S3 API parity in this specific security workflow represents a gap for organizations relying on that specific mechanism.
- **Reporting and analytics:** VAST offers exceptional SQL-queryable audit logs and analytics, but advanced active monitoring and automated reaction capabilities are currently more robust in the cloud-based uplink service than in the on-prem software. Some of these autonomous capabilities are still in the process of migrating to the local VAST AI Operating System, creating a temporary feature disparity.

PURCHASE CONSIDERATIONS

VAST Data uses a unique Gemini commercial model that decouples software licensing from hardware costs, significantly improving transparency. Customers purchase hardware at cost directly from manufacturing partners (such as Avnet) and license the software based on usable capacity. This model is transferable, allowing software licenses to move to new hardware generations without repurchasing, eliminating the hardware refresh tax.

The solution is geared toward large enterprise and CSP buyers with

petabyte-scale data needs. The architectural minimums make it unsuitable for SMBs or small-scale deployments. Deployment is flexible, supporting on-prem appliances, software-only on qualified hardware, and public cloud instances. The shared-everything architecture simplifies management by automating load balancing and failure handling, resulting in a near-zero-touch operational experience that is notably easier to manage at scale than traditional sharded clusters.

USE CASES

VAST Data supports data-intensive verticals such as life sciences, finance, media, and automotive. The platform is specifically optimized for AI/deep learning training and inference, serving as the data foundation for neoclouds and GPU-accelerated computing environments (such as NVIDIA DGX SuperPOD). It also consolidates HPC and big data analytics workloads (like from Spark or Trino) onto a single unified flash tier.

WEKA: WEKA Data Platform

SOLUTION OVERVIEW

WEKA is a data platform provider specializing in high-performance storage software designed for data-intensive workloads, particularly AI, ML, and HPC. The vendor focuses on eliminating the traditional tradeoffs between speed and simplicity by delivering a software-defined architecture that runs on commodity servers, public cloud instances, or turnkey appliances. Recently, WEKA has aggressively targeted the AI factory market, positioning its NeuralMesh technology as a converged layer that unifies file and object protocols to saturate GPUs and accelerate model training.

The WEKA Data Platform (incorporating NeuralMesh) acts as a single, unified software solution rather than a fragmented suite of modules. It is available as software-only, compatible with hardware from partners like HPE, Dell, and Supermicro, or as the WEKApod Data Platform Appliance for preconfigured deployments. The licensing model is primarily subscription-based, scaling with capacity, and includes options for cloud consumption via marketplaces.

WEKA follows a distinctive strategy that prioritizes performance and protocol versatility over traditional cheap and deep archival storage features. The platform is designed to act as a high-speed data layer where S3 is just one of several access methods (alongside POSIX and GPUDirect), allowing data ingested via object protocols to be immediately available for high-performance processing. This approach places WEKA in the Innovation hemisphere, as the solution evolves rapidly to support new AI hardware (like NVIDIA H100s) and novel data paths, differing significantly from static archival repositories.

WEKA is positioned as a Challenger and Fast Mover in the Innovation/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

WEKA scored well on a number of decision criteria, including:

- **Workload optimization:** The platform uses a zero-tuning architecture that dynamically handles mixed I/O patterns ranging from small random writes to large sequential reads without requiring manual administrative intervention. For AI pipelines, WEKA supports NVIDIA GPUDirect Storage (GDS), creating a direct path from storage NVMe to GPU memory that bypasses the CPU to reduce latency and maximize utilization during model training.
- **Storage optimization:** WEKA employs advanced data reduction techniques, including variable-block compression and global deduplication, which are optimized to function at high speeds without impeding performance. The architecture features a flash-to-object tiering mechanism that automatically moves aging data to lower-cost object storage (S3-compatible) while retaining metadata on the high-performance flash tier, ensuring directory lookups remain fast even for archived data.
- **Public cloud integration:** The solution offers native parity across on-prem and public cloud environments, with availability on AWS, Azure, GCP, and Oracle Cloud Infrastructure (OCI). The Snap-to-Object feature enables data mobility by encapsulating snapshots that can be moved between clouds or back on-prem to rehydrate clusters, facilitating true hybrid workflows in which compute follows data gravity.

OPPORTUNITIES

WEKA has room for improvement in a few decision criteria, including:

- **Auditing:** While the system generates granular logs for management and data access events, it lacks a native, integrated user interface for real-time audit analytics. Users must currently export logs to external SIEM tools like Splunk or ELK stacks to perform deep search, visualization, or forensic analysis, adding operational complexity compared to solutions with built-in audit dashboards.
- **Ransomware protection:** WEKA provides strong data resilience through immutable snapshots that can be offloaded to a locked object store bucket via Snap-to-Object. However, the S3 front end does not yet support the native S3 Object Lock API, meaning S3-native backup applications such as Veeam cannot manage immutability directly on the WEKA tier without backend configuration. The protection model is architectural rather than API-

driven, which limits interoperability with third-party data protection workflows that rely on standard Object Lock semantics.

- **Reporting and analytics:** The current WEKA Home cloud portal provides basic visibility into cluster health and capacity usage, but advanced predictive analytics are limited. A more comprehensive observability suite (NeuralMesh Observe) offering deep capacity forecasting and AIOps capabilities is now generally available for production use.

PURCHASE CONSIDERATIONS

WEKA utilizes a transparent subscription-based licensing model decoupled from the underlying hardware, which simplifies cost calculations for software-defined deployments. Pricing is generally capacity-based, and the vendor offers consumption-based options in public cloud marketplaces, preventing the sticker shock often associated with complex modular licensing. This model allows organizations to start smaller and scale costs linearly with their data growth.

The solution is engineered for performance-critical environments and is best suited for Global 2000 enterprises and specialized service providers (CSPs, MSPs) running data-intensive workloads. It is generally overengineered for SMBs seeking simple file sharing or basic backup targets. Deployment requires a specific architectural minimum to ensure resiliency, which may be a barrier for very small edge use cases, though the software-only nature allows for significant hardware flexibility.

Migration to WEKA is facilitated by its multiprotocol support, allowing legacy data to be ingested via standard NFS/SMB and immediately accessed via high-performance S3 or GDS protocols. However, organizations looking to replace a dedicated archival object store should carefully evaluate the feature gaps in S3 compliance (specifically locking and versioning) before displacing existing cheap and deep repositories.

USE CASES

WEKA primarily supports industry verticals where time to insight is critical, including life sciences (genomics), financial services (algorithmic trading), and M&E (VFX Rendering). The platform excels in AI factory use cases, for which it serves as the central data repository for the entire AI lifecycle, ingesting data via S3, processing it via GPUDirect, and archiving it via tiering. It unifies distinct storage silos (file, object, block) into a single logical namespace for diverse applications.

Zadara: Zadara Object Storage

SOLUTION OVERVIEW

Zadara differentiates itself in the market as a pioneer of the service-first model, delivering enterprise storage strictly as a fully managed STaaS offering. Rather than selling capital assets, Zadara provides on-prem and edge infrastructure that mimics the elasticity and OpEx consumption model of the public cloud while retaining the performance control of dedicated hardware. The vendor recently expanded its capabilities with the 25.07 release, introducing essential enterprise features like Unified Identity and native container-level replication.

The Zadara Object Storage solution is an integral component of the broader zStorage cloud platform. It is not a standalone silo but part of a unified architecture that delivers block, file, and object protocols on shared, multitenant infrastructure using virtual private storage arrays (VPSAs). Commercially branded as Zadara Object Storage, the offering operates within the zStorage platform and is complemented by edge computing services under the zCompute portfolio.

Zadara addresses a wide spectrum of use cases from high-performance transactional databases to massive backup repositories through its universal infrastructure. The solution is positioned in the Maturity hemisphere. The vendor prioritizes operational excellence, a 100% Uptime Guarantee, and strict AWS S3 API compatibility over the rapid release of experimental features.

Zadara is positioned as a Challenger and Fast Mover in the Maturity/Platform Play quadrant of the object storage Radar chart.

STRENGTHS

Zadara scored well on a number of decision criteria, including:

- **Auditing:** The platform generates robust, granular container logs for all object-level transactions (GET, PUT, DELETE) in standard JSON or CSV formats, enabling direct ingestion by external SIEM tools like Splunk or Azure Sentinel. The recently added Unified Identity Service enhances this capability by maintaining a consistent view of user activity and identity across distributed global regions, supporting rigorous compliance and forensics workflows.
- **Ransomware protection:** Zadara provides a defense-in-depth strategy centered on S3 Object Lock in compliance mode, which enforces a strict WORM policy that prevents data deletion by any user, including root administrators. This capability is certified Veeam Ready for Object Immutability, effectively creating a logical air gap

for backup repositories alongside the physical isolation inherent in the VPSA architecture.

- **Public cloud integration:** The vendor's cloud adjacent deployment model allows customers to place storage in colocation facilities connected to hyperscalers via low-latency links (Direct Connect, ExpressRoute). This architecture enables data to be processed simultaneously by compute resources in AWS, Azure, and GCP without incurring data egress fees, a significant economic advantage for multicloud workflows.

OPPORTUNITIES

Zadara has room for improvement in a few decision criteria, including:

- **Kubernetes support:** While the vendor offers a mature CSI driver for block and file services, it currently lacks a native CSI driver specifically for provisioning object storage buckets. Kubernetes applications must connect via standard S3 APIs or generic third-party adaptors, preventing a fully Kubernetes-native experience in which buckets and policies are managed as CRDs.
- **Workload optimization:** The solution excels at static resource isolation but lacks dynamic, policy-driven automated tiering within a single bucket. Administrators cannot set policies to automatically move aging objects between media types (such as Flash to HDD) based on access patterns; instead, they must manually provision separate pools for different performance tiers.
- **Storage optimization:** Zadara explicitly does not support inline compression or deduplication for its object storage service, features that are otherwise available on its block and file arrays. Capacity efficiency for object workloads relies entirely on erasure coding, which places the solution at a disadvantage for backup and archive use cases for which data reduction technologies could significantly lower the effective cost per gigabyte.

PURCHASE CONSIDERATIONS

Zadara utilizes a transparent, resource-based pricing model that distinguishes it from public cloud competitors. Customers pay primarily for disk capacity, with the option to add reserved engine upgrades for performance-intensive workloads, rather than being metered for API requests or total data stored. Notably, the service charges zero ingress or egress fees, eliminating the bill shock often associated with data movement in hyperscale environments.

The solution is ideally suited for MSPs and enterprises that require the data sovereignty of on-prem hardware combined with an OpEx financial model. Deployment is handled entirely as a managed service. Zadara ships the physical infrastructure to the customer's data center or edge location but retains ownership and operational responsibility. This significantly reduces deployment complexity and the burden of lifecycle management (such as firmware upgrades and drive replacements), which are performed by the vendor. However, this model means the solution is not available as software-only licensing for existing hardware.

USE CASES

Zadara supports a broad range of verticals, finding particular success with service providers (MSPs, CSPs) who white label the multitenant architecture to deliver hosted storage clouds. The platform is also strong in media and entertainment workflows requiring consistent high throughput, as well as regulated industries demanding physical data isolation. Primary use cases include immutable backup repositories (leveraging Veeam integration) and high-performance hybrid cloud storage for which data must remain accessible to multiple public clouds simultaneously without egress penalties.

A growing majority of enterprises are actively repatriating data-intensive workloads to on-prem or colocation environments, not to reject the cloud operating model but to escape the punishing economics of data gravity.

06 Analyst's Outlook

THE OBJECT STORAGE MARKET has undergone a transformation so profound that the category itself is being redefined. Once relegated to archival duty, object storage has ascended to become the primary data tier for the enterprise's most demanding workloads: generative AI, agentic automation, and high-performance computing. For IT decision-makers, the starting point for understanding this space is no longer capacity but capability. The metric that matters is not just dollars per gigabyte but object operations per second and the ability to handle the high-frequency I/O patterns generated by modern AI agents without stalling expensive GPU clusters.

Two major themes are reshaping purchase decisions: the recalibration of cloud strategies and the sovereignty imperative. The narrative of a one-way migration to the public cloud has fractured. A growing majority of enterprises are actively repatriating data-intensive workloads to on-prem or colocation environments, not to reject the cloud operating model but to escape the punishing economics of data gravity. Egress fees and API request charges have become unsustainable for AI workloads involving billions of GET and PUT requests, with a significant share of total cloud object storage spending now attributable to transaction fees rather than capacity. Simultaneously, regulatory frameworks like DORA and NIS2 are transforming data sovereignty from a policy checkbox into a rigid operational constraint, forcing enterprises to prove exactly where their data resides and who controls it.

For organizations weighing adoption, the next best action is to stop viewing object storage as a passive commodity. Audit your workloads for transactional intensity. If your data pipeline involves heavy read/write cycles (such as AI training epochs that reread data hundreds of times), standard public cloud storage classes may be economically unsustainable. Prioritize AI-native platforms that offer flat-rate economics and high-performance all-flash architectures. Look for solutions that offer a global namespace, which abstracts physical data location, allowing you to access data across edge, core, and cloud environments as a single logical pool. This capability supports data sovereignty while still leveraging cloud compute for bursty tasks. Finally, treat security as an intrinsic architectural feature. Immutable S3 Object Lock in compliance mode (where not even the root administrator can delete data before retention expires) should be mandatory in every Request for Proposal (RFP), driven by both the ransomware landscape and cyber insurance mandates.

Looking forward, the distinction between primary block storage and secondary object storage will continue to blur as UFFO architectures mature, enabling data ingested via S3 to be immediately available for high-performance processing through file interfaces without copying or gateways. Additionally, sustainability is graduating from a corporate value to a procurement mandate. With AI projected to significantly increase data center power

consumption, efficiency metrics like watts per terabyte will become as critical as performance specifications. The expansion of the Electronic Product Environmental Assessment Tool (EPEAT) ecolabel to include enterprise data storage means vendors must now demonstrate not just energy efficiency but circularity (recyclable components, conflict-free minerals, and transparent lifecycle carbon accounting).

To learn about related topics in this space, check out the following GigaOm Radar reports:

GigaOm Radar for **Primary Storage v6**

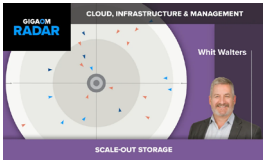

GigaOm Radar for **Scale-Out Storage v6**

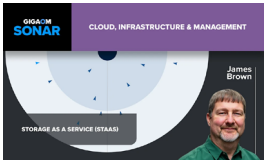

GigaOm Sonar for **Storage as a Service (STaaS) v4**

Methodology

Hitachi Vantara

*Vendors marked with an asterisk did not participate in our research process for the Radar report, and their capsules and scoring were compiled via desk research.

For more information about our research process for Radar reports, please visit our Methodology.

About Whit Walters

My mission is to deliver innovative and scalable solutions that enable data-driven decision making and business transformation. I have extensive knowledge and skills in big data, data warehousing, Apache Airflow, and Google Cloud Platform, where I hold three professional certifications. I enjoy collaborating with clients and partners, sharing best practices, and mentoring the next generation of data and cloud professionals.

GIGAOM

About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

Copyright© Knowingly, Inc. 2026 *"GigaOm Radar for Object Storage"* is a trademark of Knowingly, Inc.
For permission to reproduce this report, please contact sales@gigaom.com.

GIGAOM

GigaOm is built on the perspective of professionals who have operated, deployed, and managed the technologies we assess. Our practitioner lens brings nuance, realism, and foresight, ensuring our research is trusted, actionable, and aligned to the real pressures facing modern enterprises.