

Critical Capabilities for Object Storage

Published: 25 January 2018 **ID:** G00304492

Analyst(s): Raj Bala, Garth Landers, John McArthur

Summary

Cost reduction is the driving factor in enterprise interest in object storage, but compelling hybrid cloud storage capabilities are also attracting enterprises. We compare 13 object storage products against eight critical capabilities in use cases relevant to I&O leaders.

Overview

Key Findings

The price range for object storage platforms is wide, with startups and open-source choices pricing at the low end and more established vendors with the largest deployments pricing at the high end of the range.

Enterprises endeavor to use object APIs in their own applications, but such application modernization efforts are slow, resulting in a continued requirement for file protocols.

Hybrid cloud storage is in a renaissance period, with capabilities being reimaged with the goal of seamless operation between public cloud providers and enterprise data centers.

Analytics workloads are a large driver of data growth to public cloud object storage services, but the same is not true for on-premises object storage products due to differences in buyers and solution packaging.

Many object storage vendors claim S3 API compatibility, but testing vendor implementations reveals unfinished support, particularly as it relates to security.

Recommendations

I&O leaders responsible for accelerating infrastructure innovation and agility should:

Evaluate products from vendors that offer integrated backup appliances when considering the option to build-your-own backup solution with object storage on the back end of traditional enterprise backup products.

Identify workloads that can benefit from the elasticity provided by public cloud compute resources and shortlist object storage vendors with hybrid cloud storage capabilities that unlock such potential.

Validate vendor claims of comprehensive support for file protocols (such as NFS and SMB) through real-world performance and compatibility testing to prove the viability for their workloads.

Avoid the lowest-priced object storage product without understanding a vendor's track record, to support deployments that match their requirements.

What You Need to Know

Object storage is pervasive as the underlying platform for cloud applications that we consume in our personal lives (for example, through content streaming, photo sharing and file collaboration services). The degree of awareness and the level of adoption of object storage in the enterprise continue to be muted, but the interest is growing largely due to:

The explosion in the amount of unstructured data and the resulting need for low-cost, scalable and self-healing multitenant platforms for storing petabytes of data.

New investments in hybrid cloud storage capabilities, particularly in industries like media and entertainment, finance, and life sciences, which take advantage of the elasticity offered in the public cloud.

Growing interest from enterprise developers and DevOps team members looking for agile and programmable infrastructures that can be extended to the public cloud.

Object storage is characterized by access through RESTful interfaces that have granular, object-level security and rich metadata that can be tagged to it. Object storage products are available in a variety of deployment models — virtual appliances, managed hosting, purpose-built hardware appliances or software that can be installed on standard server hardware. These products are capable of huge scale in capacity, and many of the vendors included in this research have production deployments beyond 10PB. On-premises object storage is designed for workloads that require high bandwidth and is rarely used for transactional workloads that demand high input/output operations per second (IOPS) and low latency.

The new generation of object storage products relies mainly on erasure-coding schemes that can improve availability at lower-capacity overhead and cost, when compared with the traditional redundant array of independent disks (RAID) schemes. The overwhelming support for the Amazon Simple Storage Service (S3) application programming interface (API) among the object storage vendors is stimulating market demand for these products, although the level of compatibility with the S3 API varies widely. And there can still be lock-in due to proprietary methods of managing metadata.

In the vendor sections that follow the scoring, we describe three critical traits of each product (metadata implementation, support for file protocols and hybrid cloud storage capabilities):

1. We note a vendor's metadata implementation because this is often one of the most challenging components to scale in a distributed system, and the differences between products in this market vary significantly.

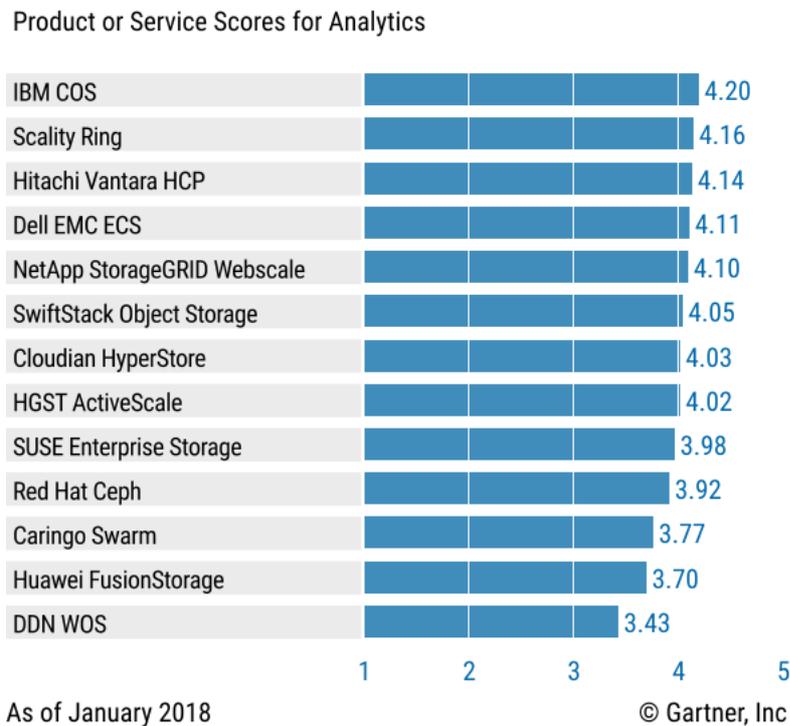
- Enterprises continue to require file protocols (such as NFS and SMB) in their environments. This has pushed object storage vendors to include such protocols, but support varies between no support at all to native, platform-level capabilities. We note these capabilities particularly when implemented through a gateway that presents a single point of failure compared to a native, horizontally scaled protocol with no single point of failure.
- We describe a vendor's hybrid cloud storage capabilities because of the early customer adoption utilizing storage products that act as a bridge between on-premises and public cloud environments. Additionally, significant differences exist between vendors as they relate to hybrid cloud capabilities.

IT leaders who need highly scalable, self-healing and cost-effective storage platforms for large amounts of unstructured data should evaluate the suitability of object storage platforms. They should use this research as a basis to identify the appropriate products for their use cases.

Analysis

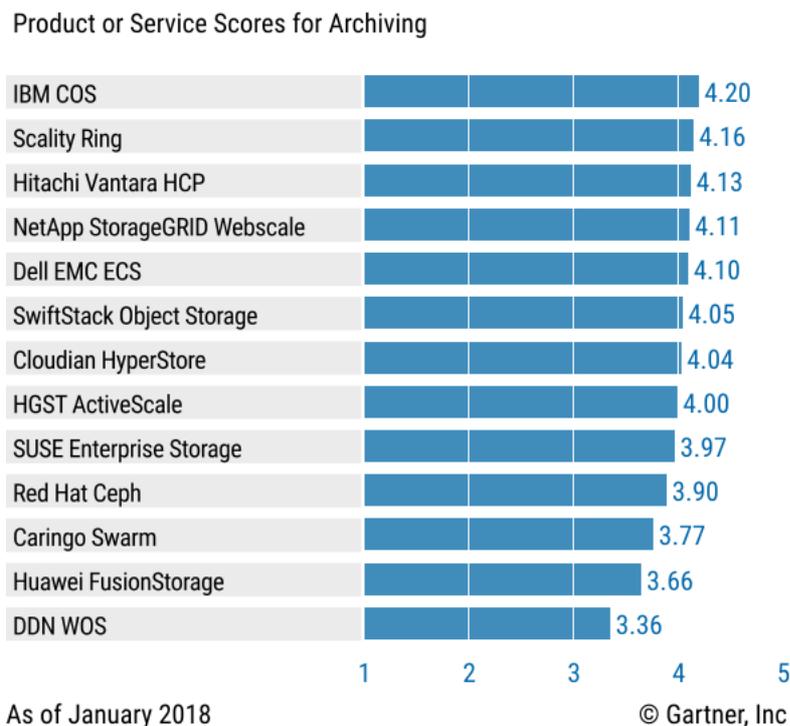
Critical Capabilities Use-Case Graphics

Figure 1. Vendors' Product Scores for the Analytics Use Case



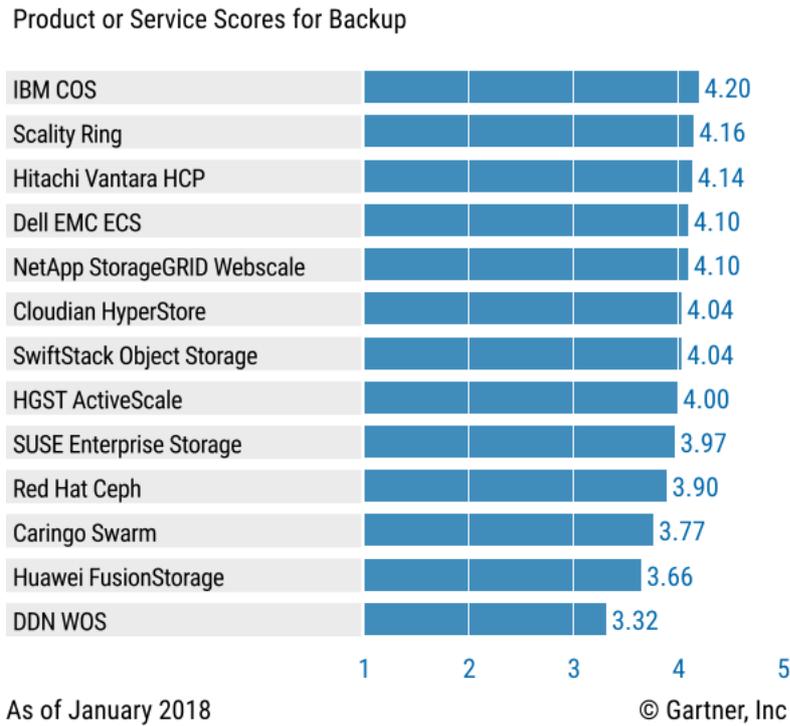
Source: Gartner (January 2018)

Figure 2. Vendors' Product Scores for the Archiving Use Case



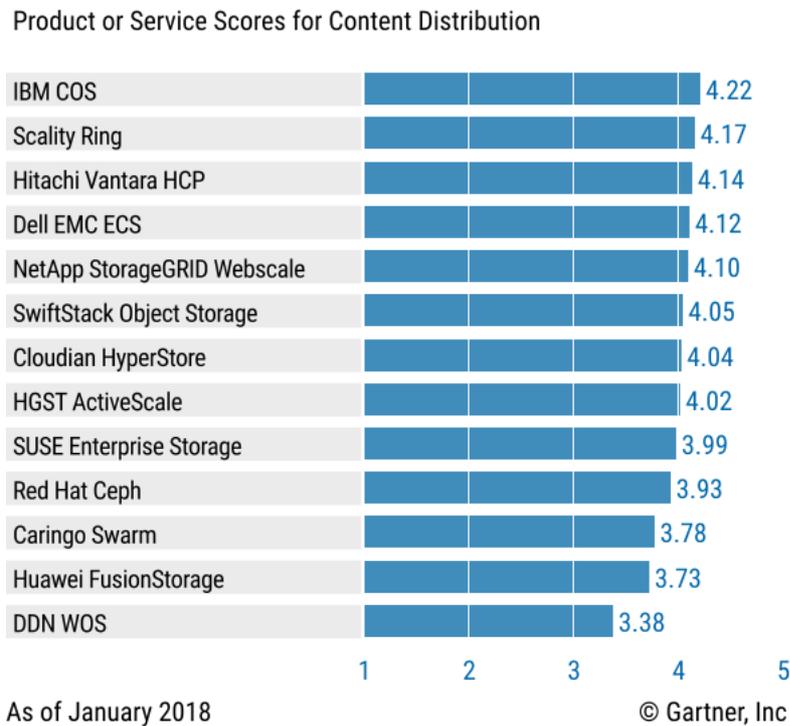
Source: Gartner (January 2018)

Figure 3. Vendors' Product Scores for the Backup Use Case



Source: Gartner (January 2018)

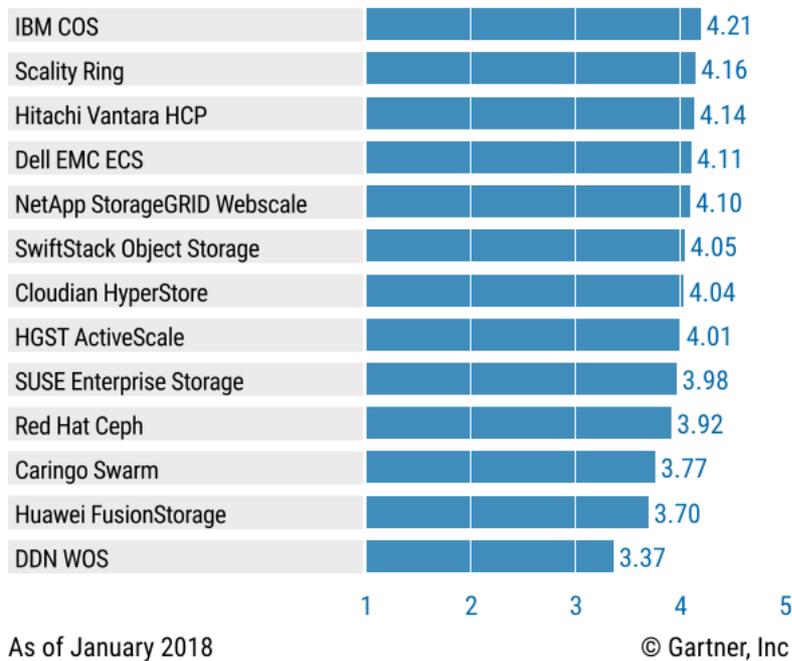
Figure 4. Vendors' Product Scores for the Content Distribution Use Case



Source: Gartner (January 2018)

Figure 5. Vendors' Product Scores for the Cloud Storage Use Case

Product or Service Scores for Cloud Storage



Source: Gartner (January 2018)

Vendors

Caringo Swarm

Swarm is an object storage platform by Caringo that is often paired with FileFly, a solution for tiering content originating in Windows and NetApp filers. Swarm is a mature product developed by a team with a long track record in engineering object storage solutions. Swarm's access control settings are widespread with support for cluster, tenant, domain, bucket and individual objects. Swarm is a good fit for governance-oriented archiving, as the product has support for write once, read many (WORM), content authenticity and legal hold capabilities. Deployment can occur on bare-metal servers in addition to virtual machines with support for vSphere and Hyper-V. Caringo is particularly strong in the healthcare vertical.

Caringo implements Amazon S3 API compatibility, but does not support Amazon Web Services (AWS) as a target in hybrid cloud scenarios. The Swarm management dashboard is rather basic and not as feature-rich as competitive alternatives lacking strong reporting capabilities.

Notable Product Traits:

Metadata: Metadata storage in-line with objects.

File Protocols: NFSv4 implemented as native, platform-level protocol. SMB available through FileFly, an OEM of Moonwalk's data management product.

Hybrid Cloud Storage: Simple tiering of data to Microsoft Azure blob storage.

Cloudian HyperStore

Cloudian develops HyperStore, an object storage platform deployed on-premises that is significantly influenced by Amazon S3, both in terms of API support and portal aesthetics. As such, enterprises seeking a platform with usability and features akin to S3 will feel comfortable using HyperStore. HyperStore has a well-designed dashboard for viewing historical capacity consumption, making it readily noticeable when capacity is decreasing with increased usage. The same at-a-glance dashboard provides information on current throughput and transactions to the cluster. Such functionality is typically difficult to ascertain in most other object storage products. HyperStore also supports AWS-style IAM policies to manage access to objects.

HyperStore sends an incorrect response for a number of requests to its platform when tested for "bug compatibility" with the Amazon S3 API. This could lead to inconsistencies in application behavior from apps that expect certain compatibility with tools in the Amazon S3 ecosystem.

HyperStore's NFS and SMB capabilities are implemented as a cloud storage gateway rather than native protocols on the platform itself.

Notable Product Traits:

Metadata: Metadata storage using Apache Cassandra.

File Protocols: No native support for NFS or SMB.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

DDN WOS

Shipping for over five years, WOS is an object storage product originally designed for content distribution, but the primary use case for WOS today is archiving of high-performance computing data. WOS is integrated with DDN's GRIDScaler and EXAScaler parallel file systems appliances and space-efficiently supports rapid access to billions of objects and hundreds of petabytes of storage. WOS has an API for integration with management applications, as well as an optional S3 and Swift interface for access to third-party applications. For data resiliency, in addition to replication, WOS supports multiple levels of erasure coding at the system, site and multisite level.

WOS's S3 implementation uses Apache Hbase as an external metadata store that is separate from WOS's native metadata store. Effectively, DDN has written a new object store around their existing object store as a means of quickly getting to market with S3 compatibility. Such inelegant design decisions result in the product's inability to scale well beyond single-site deployments. Additionally, WOS's S3 implementation does not support server-side encryption.

Notable Product Traits:

Metadata: Metadata storage for DDN's S3 API using Apache Hbase. Metadata storage for the native WOS API stored in-line with objects.

File Protocols: NFS and SMB offered through Lustre or the WOS bridge.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

Dell EMC ECS

Elastic Cloud Storage (ECS) is Dell EMC's third-generation object storage platform. Much of the lessons learned from the company's previous generations are readily apparent in ECS. As a result, ECS has modern aesthetics both from the perspective of an elegantly designed user interface and the product's underlying architecture. ECS is based on a highly scalable, purpose-built key/value store that serves as the platform's metadata store. The ECS monitoring portion of the user interface has useful information, such as utilization and capacity metrics, in addition to the health status of portions of the cluster. ECS supports a "transformation" process that allows for a gradual, seamless migration from Centera, the company's widely deployed object storage product focused on governance and compliance.

ECS lacks fine-grained controls to segregate object storage users, such as setting granular permissions in terms of access to objects by API functions. ECS's NFS implementation is not optimized for general-purpose enterprise file workloads like other Dell EMC products (such as Isilon).

Notable Product Traits:

Metadata: Metadata storage using a company-developed, distributed key/value store.

File Protocols: NFSv4 implemented as native, platform-level protocol.

Hybrid Cloud Storage: Deployment includes private cloud instantiations in Virtustream data centers.

HGST ActiveScale

ActiveScale is the most recent in the evolution of HGST's object storage products. ActiveScale has a well-designed dashboard with useful at-a-glance metrics, such as the percentage of capacity used, data durability state and system performance characteristics. ActiveScale also allows administrators to delve deeper into each of these categories for a finer-grain view of specific measures as needed. ActiveScale provides a unique view into the underlying hardware, such as the state of SSDs, HDDs, network cards and power supplies. Additionally, ActiveScale engineers have contributed to the Hadoop S3A plug-in, resulting in depth-of-knowledge about big data workloads not found in most other object storage vendors.

ActiveScale is a substantially rewritten product compared with the company's object storage products that came before it. The immaturity of ActiveScale is apparent when viewed through a security lens, as the product is missing features, such as object-level access control lists (ACLs) and restricting functional access by user. Furthermore, ActiveScale fails a wide array of tests that exercise its S3 API compatibility. Upgrading from ActiveArchive, the previous HGST object storage product, to ActiveScale requires a disruptive, forklift upgrade.

Notable Product Traits:

Metadata: Distributed key/value store using LMDB, a B-tree-style database.

File Protocols: No native support for NFS or SMB.

Hybrid Cloud Storage: No notable support for hybrid cloud storage.

Hitachi Vantara HCP

Hitachi Vantara has four offerings in its portfolio, with Hitachi Content Platform (HCP) at the center as the object storage solution. HCP Anywhere is deployed as a content collaboration platform. Hitachi Data Ingestor (HDI), a cloud file gateway, and Hitachi Content Intelligence provide search and analytics insights. HCP has a large and varied community of third-party ISV support, particularly for archiving. In relation to that, the product supports S3 and can act as an intermediary for applications like Splunk, Alfresco, Dragondisk and others without requiring HCP-specific integration or customization. HCP is also frequently chosen based on attractive pricing.

Some Hitachi HCP customers have cited customer support as an area in need of improvement, specifically with product expertise and the time needed to resolve issues. The product supports delivery as either software-only running in a virtual machine or a physical appliance, but is not available as a bare-metal deployment. Each deployment of HCP is autonomous meaning there isn't a unified approach for management of multisite environments. While the management dashboard is feature-rich, it can be hard for users to navigate, and the layout could be improved.

Notable Product Traits:

Metadata: Metadata storage using a company-developed, distributed key/value store.

File Protocols: NFSv4 and SMBv3 implemented as native, platform-level protocols.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

Huawei FusionStorage

FusionStorage is a new product, released as the successor to Huawei's previous object storage forays, which included OceanStor UDS.

FusionStorage is a distributed storage system with block, file and object storage capabilities. FusionStorage is deployed on-premises in enterprise data centers, but also serves as the underlying storage layer for Huawei's public cloud offerings, which include enablement of third-party providers, including China Telecom, Deutsche Telekom, Orange and Telefonica. The FusionStorage administrator interface is a single pane of glass to manage all storage interfaces, including object storage APIs. Huawei positions the object storage capabilities of FusionStorage for archiving and backup use cases. As such, FusionStorage has support for backup products from vendors, such as Commvault and Veritas.

Despite Huawei assigning a "6.0" version to the product, FusionStorage 6.0 is the first version of the platform that integrates block, file and object API support. The product lacks key features, such as compression and distributed erasure coding to support durability across geographies.

FusionStorage does not have fine-grained security controls that allow restricting API functions on data by users, as is becoming more frequent. The

usability of the FusionStorage's web interface for managing object storage significantly trails its competitors.

Notable Product Traits:

Metadata: Metadata storage using a peer-to-peer architecture.

File Protocols: NFSv3, NFSv4, SMBv2 and SMBv3 implemented as native, platform-level protocols.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

IBM COS

Cloud Object Storage (COS) serves as the underlying infrastructure for both on-premises and public cloud deployments in reference to IBM's object storage offerings. In the context of this research note, we are only evaluating the on-premises version of COS. COS has configurable erasure coding, object-level ACLs, built-in quota management and recently introduced WORM support through proprietary extensions to COS's implementation of the Amazon S3 API. WORM capabilities are achieved through COS's erasure coding algorithms, making it more efficient compared to alternative approaches that rely on the replication of complete objects to achieve compliance. Prospective COS customers previously required a minimum commitment of 288TB, but is now available at much smaller configurations starting at 72TB.

COS does not have fine-grained security controls in the context of defining access to objects by API functions. Developers accustomed to interacting with Amazon S3 will find COS's user interface to be markedly less intuitive for common tasks, such as generating the necessary information to begin writing to the nodes. Bulk delete operations can render all data in COS to be completely inaccessible due to exhausted RAM resources. As a result, customers will need to work with IBM support to tune the delete performance to match their specific workloads to mitigate this issue.

Notable Product Traits:

Metadata: Company-developed, distributed key/value store.

File Protocols: Limited NFSv3 support currently lacking support for features, such as hard links. Requires use of an external data store used rather than the native metadata store.

Hybrid Cloud Storage: Deployment options include private cloud instantiations in IBM data centers.

NetApp StorageGRID Webscale

NetApp offers StorageGRID Webscale, an object storage platform available as a virtual or physical appliance. With extensive support for ISVs, StorageGRID is well-positioned in the archiving and backup use cases when combined with NetApp's AltaVault, a backup-focused cloud storage gateway. The company has deep engineering relationships with public cloud providers (such as AWS and Microsoft) and is making significant innovations with respect to hybrid cloud deployments as a result. Similarly, StorageGRID endeavors to provide a public-cloud-like experience to its on-premises deployments, and thusly provides an AWS-like IAM experience to governing access to data. StorageGRID also integrates with Elastic Search to allow metadata to be searchable.

StorageGRID's current hybrid cloud integration does not support bidirectional synchronization to public cloud targets, as it lacks the necessary integration to manage data changed in the public cloud. StorageGRID's NFS and SMB support are limited by the company's gateway implementation rather than protocols supported on the distributed platform itself.

Notable Product Traits:

Metadata: Metadata storage using Apache Cassandra.

File Protocols: No native support for NFS or SMB; relies on company-developed gateway.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

Red Hat Ceph

Red Hat is the primary developer behind Ceph, an open-source object storage solution delivering block, file and object storage capabilities. The primary use cases are for OpenStack-based block and object, big data, content distribution networks, and active archives. Ceph supports unlimited nodes and capacity, and automatically rebalances and scales when adding additional servers to the pool. Ceph offers three levels of erasure coding with a maximum usable capacity of 73% of raw capacity. Additional erasure coding schemes are enabled through plug-ins. Ceph offers deduplication and compression through recently acquired Permabit. Management is through CLI, REST APIs and a GUI.

Reference customers strongly recommend that users leverage Red Hat's storage consulting group to properly size the cluster for performance and capacity, and conduct thorough testing at scale and under failure scenarios. Performance tiering is only available within a node. Customers expressed disappointment with the current maturity of GUI-based management.

Notable Product Traits:

Metadata: All metadata stored in-line with objects.

File Protocols: No native support for SMB. Native support for NFSv3 and NFSv4.

Hybrid Cloud Storage: No notable support for hybrid cloud storage.

Scality Ring

Scality develops Ring, a distributed storage platform with native support for both file protocols and object storage APIs. Platform-level support for these primary access methods distinguishes Scality as being one of the few vendors that support multiprotocol access to the same objects. This is useful in scenarios where objects are ingested with file protocols and consumed with the S3 API. Scality supports WORM and compliance to support industries with regulations that govern the storage of data. Scality has developed a container-deployable runtime version of its S3 implementation to allow developers to have a Scality experience without needing to deploy hardware in a data center.

Scality promotes a multicloud architecture with its Zenko product, but such application architectures are very rare due to the difficulty in managing the complexities of supporting multiple providers within one application's architecture. Among comparable offerings in this report, Scality Ring was among the vendors with the longest deployment times. Further, reference customers have cited that, while ultimately happy with the product, Scality field service personnel were often required to deploy additional clusters due to challenging deployment requirements.

Notable Product Traits:

Metadata: Metadata storage using LevelDB.

File Protocols: NFSv3, SMBv2 and SMBv3 implemented as native, platform-level protocols.

Hybrid Cloud Storage: Simple tiering of data to S3-compatible targets.

SUSE Enterprise Storage

SUSE Enterprise Storage is based on the open-source Ceph project that delivers block, file and object storage services. The product can be purchased as a software-only solution, together with servers as a reference architecture or as an appliance from OEM partners. Pricing is for support and is storage-node-based, with entry solutions starting at four storage nodes and 100TB. Pricing includes six infrastructure nodes for management, monitoring, metadata and file access gateways. Common uses include backup, archiving and bulk storage. The solution supports automated tiering and object promotion and demotion, based upon user-definable parameters, and supports encryption of data in flight and at rest. SUSE provides management via CLI, web interface and REST APIs.

SUSE Enterprise Storage does not currently support data deduplication. The solution may be a poor fit for small objects, as the system does not self-tune based on object size. Reference customers expressed concern regarding the maturity of the GUI. As a highly configurable, non-self-tuning product, SUSE Enterprise Storage requires more training and expertise to get from installation to optimization. Pre- and post-installation consulting days may be required to ensure that the solution is properly tuned for customer-specific workloads.

Notable Product Traits:

Metadata: All metadata stored in-line with objects.

File Protocols: No native support for SMB. Native support for NFSv3 and NFSv4 (tech preview).

Hybrid Cloud Storage: No notable support for hybrid cloud storage.

SwiftStack Object Storage

SwiftStack is the largest contributor to Swift, the open-source object storage component in the OpenStack suite of products. Though SwiftStack has its roots in open source, it has built a number of proprietary enterprise capabilities around Swift, such as cluster management and S3 API compatibility. SwiftStack has also built unique capabilities that allow a customer to deploy object storage clusters on-premises, but administer them using a SwiftStack-hosted and managed portal in the cloud. SwiftStack is one of the earliest object storage vendors to support modern forms of hybrid cloud storage that enable enterprises to synchronize data to the public cloud to use its elastically scale compute resources.

SwiftStack fails a number of basic tests of the product's S3 API that it should otherwise handle appropriately. Further, the product fails or otherwise does not implement many aspects of the S3 API particularly focused on versioning. SwiftStack has limited support for SMBv3.

Notable Product Traits:

Metadata: Metadata storage using file-system-extended attributes or alongside object itself.

File Protocols: NFSv3, SMBv1, SMBv2 and SMBv3 (limited) implemented as native, platform-level protocols.

Hybrid Cloud Storage: Bidirectional synchronization of data to S3-compatible targets.

Context

The first generation of object storage, in early 2000, manifested as content-addressed storage (CAS). During the late 2000s, the second phase of object storage shifted the product focus to cloud uses, with a development emphasis on a cost-effective cloud storage infrastructure, with erasure codes for storage-efficient protection and better WAN support.

Cloud computing, as a category, was started in 2006 with the introduction of Amazon S3, an object storage service. So it should be no surprise that the current market for object storage is largely being driven by S3's originator, AWS.

There's a symbiotic relationship between Amazon S3 and other object storage services and products. AWS innovations and use cases on S3 guide the market for vendors offering both public-cloud-based and on-premises object storage.

However, there's a delay between when other object storage vendors implement new functionality and when the use cases materialize. In many instances, the public cloud use cases of object storage never materialize in the enterprise data center because of differences between the surrounding set of solutions and the buyers in each segment. This often results in Mode 1 workloads using on-premises object storage and Mode 2 workloads in the public cloud.

Product/Service Class Definition

Object storage refers to storage hardware and software infrastructure that house data in structures called "objects," and serve hosts via protocols (such as HTTP) and APIs (such as Amazon Simple Storage Service [Amazon S3]). Conceptually, objects are similar to files, in that they are composed of content and metadata. In general, objects support richer metadata than file storage by enabling users or applications to assign attributes to objects that can be used for administrative purposes, data mining and information management.

Critical Capabilities Definition

Object storage products often outscore traditional block and file storage products in capacity scalability, security/multitenancy, total cost of ownership (TCO) and manageability, although they tend to lag in performance, interoperability and efficiency. Given the nascent state of the market, several features that clients expect in a traditional NAS system may be absent or less developed in object storage products, due to design considerations or product immaturity. Clients need to understand these trade-offs during the procurement process.

Enterprises should consider the following eight critical capabilities when deploying object storage products. Enterprises can work toward these goals by evaluating object storage products in all capability areas.

Capacity

The ability of the product to support growth in capacity in a nearly linear manner. It examines object storage capacity scalability limitations in theoretical and real-world configurations, such as maximum theoretical capacity, maximum object size and maximum production deployment.

Storage Efficiency

The ability of a product to support storage efficiency technologies, such as compression, single-instance storage/deduplication, tiering, erasure coding and massive array of idle disks (MAID) to reduce TCO.

Interoperability

The ability of the product to support multiple networking topologies, third-party ISV applications, public cloud APIs and various deployment models.

Manageability

The automation, management, monitoring, and reporting tools and programs supported by the product. In addition, ease of setup and configuration and metadata management capabilities were considered.

These tools and programs can include single-pane management consoles, monitoring systems and reporting tools designed to help personnel seamlessly manage systems, monitor system usage and efficiencies, and anticipate and correct system alarms and fault conditions before or soon after they occur.

Performance

The per-node and aggregated throughput for reads and writes that can be delivered by the cluster in real-world configurations.

Resilience

The platform capabilities for providing high system availability and uptime. Options include high tolerance for simultaneous disk and/or node failures, fault isolation techniques, built-in protection against data corruption, and data protection techniques, such as erasure coding and replication.

Features are designed to meet users' recovery point objectives (RPOs) and recovery time objectives (RTOs). There are several methods for data protection in today's object storage products. RAID is becoming less popular due to huge capacity overheads and long rebuild times. The simplest way to protect data is replication, which stores multiple copies of the data locally or in a distributed manner. A more innovative data protection scheme is erasure coding, which breaks up data into "n" fragments and "m" additional fragments across n+m nodes, offering clients configurable choices, depending on their cost and data protection requirements. Enterprises often combine erasure coding and replication, because the former performs well with large files, whereas the latter works well when there are large numbers of small files. WAN costs and performance considerations in distributed environments are also factors.

Security and Multitenancy

The native security features embedded in the platform provide granular access control, enable enterprises to encrypt information, provide robust multitenancy, offer data immutability and ensure compliance with regulatory requirements.

Value

The price of the product relative to the capabilities an enterprise stands to experience.

Use Cases

Analytics

This applies to storage consumed by big data analytics applications.

Interoperability, performance (more specifically, bandwidth), resilience and scalability are critical to success. These include features to tolerate disk/node failures, versioning to facilitate check-pointing of long-running applications and bandwidth to reduce time to insight.

Archiving

The earliest enterprise use case for object storage products, it has been used for more than a decade. Products provide cost-effective, scalable, long-term data preservation.

For this use case, object storage products are used to store immutable data for decades. Features such as WORM, legal hold and object-level versioning increase the attractiveness of object storage as an archiving target in terms of ease of access, affordability and data immutability. Security, resilience, interoperability and manageability (such as indexing and metadata management features) are important selection considerations, and are heavily weighted.

Backup

I&O leaders have used object storage products as backup targets for years because they provide added scalability for large backup datasets.

Object storage is important for meeting increasing demands for disk-based backup at lower cost. Resilience, storage efficiency, performance and interoperability with a variety of backup ISVs are important selection considerations, and are heavily weighted.

Cloud Storage

This is the most prominent use case for object storage products. Most popular consumer and enterprise public clouds are built on an object storage foundation.

This use case refers broadly to service-provider-built public and private clouds and enterprise-built public, community, hybrid and private clouds, where the access is through HTTP. This is different from VM storage (which is likely to be block storage) that providers or enterprises build for high-performance applications, such as databases. Resilience, capacity, performance and security are the most important considerations in the choice of products, and are heavily weighted.

Content Distribution

This refers to distributed delivery of content for users across multiple locations to enhance collaboration and distribution of data.

Intelligent and predictive content placement that is served via optimal network routes with high availability, high performance and robust data integrity are key considerations. Performance, resilience, interoperability and manageability are critical selection considerations, and are heavily weighted.

Vendors Added and Dropped

Added

SUSE qualified for this Critical Capabilities since its last publication and was added as a result.

Dropped

No vendors were dropped from this research note since its last publication.

Inclusion Criteria

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

Earn revenue above \$10 million per year for the object storage product between 1 May 2016 and 30 April 2017, or have at least 50 production customers each consuming more than 300TB capacity through object storage protocols only. Vendor must provide reference materials to support this criterion.

The product must be installed in at least two major geographies (among North America, EMEA, Asia/Pacific and South America).

The product should be deployed across multiple use cases that are outlined in Critical Capabilities for Object Storage.

The product must be designed for primarily on-premises workloads and not as a pass-through solution where data will be permanently stored elsewhere.

The vendor should own the storage software intellectual property and be a product developer. If the product is built on top of open-source software, the vendor must be one of the top 10 active contributors to the community (in terms of code contribution).

Packaging:

The product must be sold as either an appliance or a software-based solution.

The product must be available for purchase as a stand-alone storage product rather than an integrated or converged system with compute and hypervisor bundle.

Product capabilities:

The product must have object access to the common name space.

The product must have a shared nothing architecture where data is replicated or erasure coded over the network across multiple nodes in the cluster. It must have the ability to handle disk, enclosure or node failures in a graceful manner without impacting availability.

The global namespace must be capable of 1PB expansion.

The cluster must span more than three nodes.

Support for horizontal scaling of capacity and throughput in a cluster mode or independent node additions with a global namespace.

Table 1. Weighting for Critical Capabilities in Use Cases

Critical Capabilities	
Value	
Analytics	13%
Archiving	15%
Backup	15%
Content Distribution	12%
Cloud Storage	13%
Storage Efficiency	
Analytics	12%
Archiving	14%
Backup	10%

Content Distribution	14%
Cloud Storage	12%
Security and Multitenancy	
Analytics	9%
Archiving	13%
Backup	14%
Content Distribution	13%
Cloud Storage	13%
Capacity	
Analytics	12%
Archiving	13%
Backup	14%
Content Distribution	13%
Cloud Storage	12%
Interoperability	
Analytics	16%
Archiving	13%
Backup	14%
Content Distribution	13%
Cloud Storage	13%
Manageability	
Analytics	13%
Archiving	13%
Backup	11%
Content Distribution	10%

Cloud Storage	13%
Performance	
Analytics	13%
Archiving	11%
Backup	11%
Content Distribution	12%
Cloud Storage	12%
Resilience	
Analytics	12%
Archiving	8%
Backup	11%
Content Distribution	13%
Cloud Storage	12%
Total	
Analytics	100%
Archiving	100%
Backup	100%
Content Distribution	100%
Cloud Storage	100%

Source: Gartner (January 2018)

This methodology requires analysts to identify the critical capabilities for a class of products/services. Each capability is then weighed in terms of its relative importance for specific product/service use cases.

Critical Capabilities Rating

Table 2. Product/Service Rating on Critical Capabilities

Critical Capabilities	
Value	
Caringo Swarm	3.2
Cloudian HyperStore	3.9
DDN WOS	2.3

Dell EMC ECS	3.5
HGST ActiveScale	3.6
Hitachi Vantara HCP	4.0
Huawei FusionStorage	2.0
IBM COS	3.8
NetApp StorageGRID Webscale	4.0
Red Hat Ceph	3.0
Scality Ring	4.0
SUSE Enterprise Storage	3.5
SwiftStack Object Storage	3.8
Storage Efficiency	
Caringo Swarm	3.8
Cloudian HyperStore	4.0
DDN WOS	4.1
Dell EMC ECS	4.2
HGST ActiveScale	4.1
Hitachi Vantara HCP	4.0
Huawei FusionStorage	4.1
IBM COS	4.4
NetApp StorageGRID Webscale	4.2
Red Hat Ceph	4.2
Scality Ring	4.2
SUSE Enterprise Storage	4.2
SwiftStack Object Storage	4.2
Security and Multitenancy	
Caringo Swarm	4.0
Cloudian HyperStore	4.3
DDN WOS	2.0
Dell EMC ECS	4.2
HGST ActiveScale	3.8
Hitachi Vantara HCP	4.1
Huawei FusionStorage	3.9

IBM COS	4.3
NetApp StorageGRID Webscale	4.1
Red Hat Ceph	4.0
Scality Ring	4.2
SUSE Enterprise Storage	4.0
SwiftStack Object Storage	3.9
Capacity	
Caringo Swarm	4.0
Cloudian HyperStore	4.1
DDN WOS	4.0
Dell EMC ECS	4.3
HGST ActiveScale	4.2
Hitachi Vantara HCP	4.2
Huawei FusionStorage	3.9
IBM COS	4.4
NetApp StorageGRID Webscale	4.2
Red Hat Ceph	4.1
Scality Ring	4.2
SUSE Enterprise Storage	4.1
SwiftStack Object Storage	4.2
Interoperability	
Caringo Swarm	4.0
Cloudian HyperStore	4.0
DDN WOS	3.5
Dell EMC ECS	4.3
HGST ActiveScale	4.2
Hitachi Vantara HCP	4.2
Huawei FusionStorage	4.0
IBM COS	4.2
NetApp StorageGRID Webscale	4.1
Red Hat Ceph	4.0
Scality Ring	4.2

SUSE Enterprise Storage	4.0
SwiftStack Object Storage	4.1
Manageability	
Caringo Swarm	3.8
Cloudian HyperStore	4.2
DDN WOS	4.1
Dell EMC ECS	4.2
HGST ActiveScale	4.1
Hitachi Vantara HCP	4.2
Huawei FusionStorage	3.7
IBM COS	4.2
NetApp StorageGRID Webscale	4.2
Red Hat Ceph	4.0
Scality Ring	4.1
SUSE Enterprise Storage	4.0
SwiftStack Object Storage	4.2
Performance	
Caringo Swarm	3.7
Cloudian HyperStore	3.8
DDN WOS	3.6
Dell EMC ECS	4.0
HGST ActiveScale	4.0
Hitachi Vantara HCP	4.2
Huawei FusionStorage	4.0
IBM COS	4.1
NetApp StorageGRID Webscale	4.0
Red Hat Ceph	4.0
Scality Ring	4.2
SUSE Enterprise Storage	4.0
SwiftStack Object Storage	3.9
Resilience	
Caringo Swarm	3.7