IDC MarketScape

IDC MarketScape: Worldwide Object-Based Storage 2018 Vendor Assessment

Amita Potnis

THIS IDC MARKETSCAPE EXCERPT FEATURES NETAPP

IDC MARKETSCAPE FIGURE

FIGURE 1

IDC MarketScape Worldwide Object-Based Storage Vendor Assessment

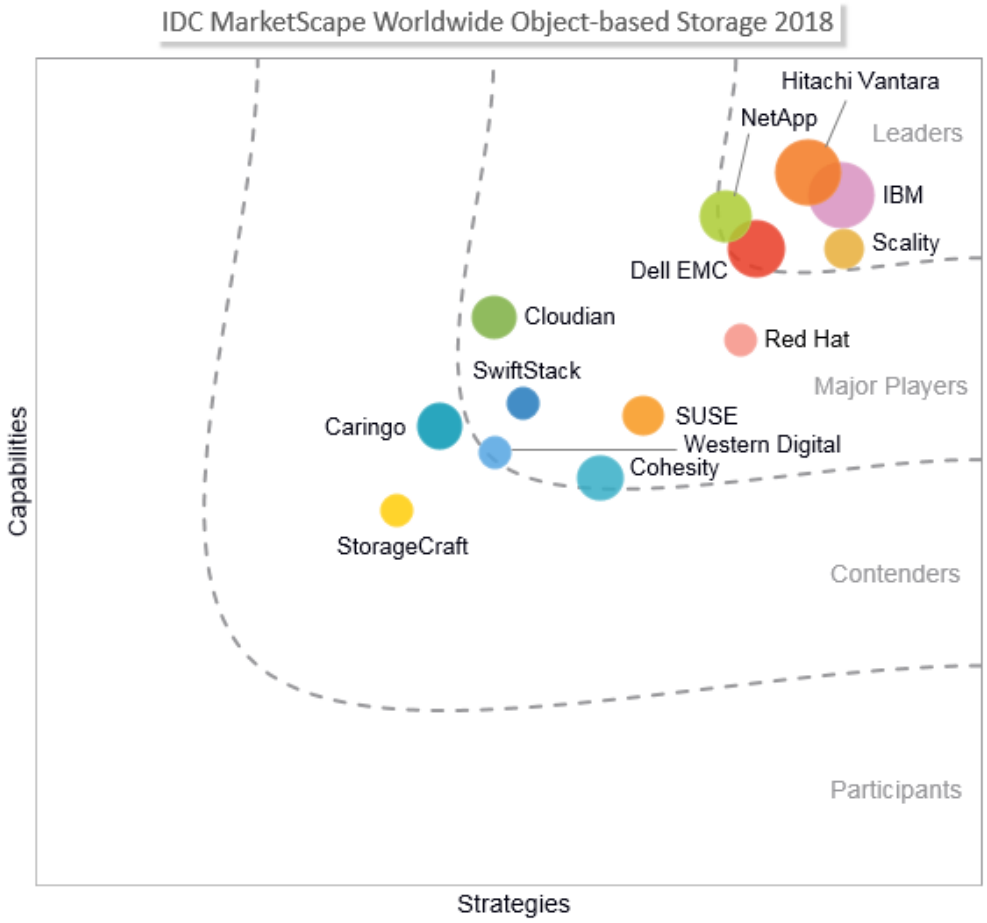

Source: IDC, 2018

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

June 2018, IDC #US42665518e

IDC OPINION

In *IDC MarketScape: Worldwide Object-Based Storage 2016 Vendor Assessment* (IDC #US41918416, December 2016), we stated that the problems of storing and managing unstructured data are not going away. This remains true in 2018. According to *File- and Object-Based Storage Survey Findings, 2017: Adoption and Workloads Trends — Part 1* (IDC #US43630018, March 2018), more than 50% of the 450 respondents adopted object-based storage (OBS) because existing NAS was too complex to manage. More than 40% of the respondents also indicated that scalability was a major concern for their organizations given the ever-growing data sets. Organizations large and small are now aware of the advantages of OBS and the flexibility of offerings of this technology. As adoption of OBS continues to increase, much of the feature functions are now standard across products. Examples of such features are S3 compatibility, erasure coding, replication, and custom metadata. IDC believes that the OBS market segment has reached a maturity point, and going forward the strategic vision in terms of partnerships with public cloud providers, file system capabilities, a portfolio of diverse solutions that support a variety of use cases across verticals, and so forth will play a strategic role in market penetration. Based on several customer interviews, IDC expects OBS will take a more prominent role in production workloads and this is where the OBS suppliers will see growth.

In this IDC MarketScape, IDC assesses the present commercial OBS supplier (suppliers that deliver software-defined OBS solutions as software or appliances much like other storage platforms) landscape. Cloud-based storage services based on OBS are not included in this IDC MarketScape. Since the publication of the last OBS IDC MarketScape, the landscape has changed dramatically — a fragmented OBS market has now somewhat consolidated. As demand for traditional external storage systems continues to decline, several mainstream storage suppliers enhanced their portfolio to include OBS solutions by developing them in-house or via acquisitions. This IDC MarketScape assesses 13 OBS suppliers that are "owners of intellectual property (IP)."

IDC believes that going forward, the longevity and success of any OBS supplier lie in its strategy and road map. Today, there is a definitive shift to compete together as opposed to commercial OBS suppliers competing against public cloud. In recent months, several OBS suppliers (listed in this IDC MarketScape) have partnered with public cloud storage service providers to support end users' multicloud initiatives as well as a wider range of workloads. Therefore, IDC would like its readers to note that this IDC MarketScape gives additional weight to strategy than capability. In earlier IDC MarketScapes, strategy and capability criteria carried equal weightage. When considering an OBS solution, IDC recommends that end users look for the following capabilities and business strategies of OBS suppliers that are representative of the market:

- **Security:** What data would move to the OBS platform, and what kind of security services does the potential product support (encryption, user authentication, protection against ransomware, unforeseen events, etc.)?
- **Multicloud strategy initiatives:** Does the commercial OBS solution support your organization's multicloud initiative by supporting data tiering to public cloud? Is data moved in native format supported by the public cloud, or does it need to be accessed through the OBS solution's proprietary interface?
- **Performance:** Will the OBS solution support higher performance needs for specific production workloads? More specifically, does it support file access protocols (NFS and SMB) natively via a file system?

- **Information life-cycle management (ILM):** Does the OBS solution support a robust set of ILM policies — not just for on-premises deployments but also in the public cloud/tape or optical media?
- **Cost:** Does the OBS solution lend itself to lower capex or opex, and will it sustain cost savings in a 5- to 10-year period?

IDC MARKETSCAPE VENDOR INCLUSION CRITERIA

This IDC study assesses the capabilities and business strategies of key suppliers in the (scale-out) OBS market segment — which is part of the overall file- and object-based storage market. This evaluation is based on a comprehensive framework and a set of parameters that gauge the success of a supplier in delivering an OBS solution in the market. This study includes analysis of 13 notable players in the commercial OBS market, with broader portfolios and global scale. The suppliers enlisted in this study are (in alphabetical order) Caringo, Cloudian Inc., Cohesity, Dell EMC, Hitachi Data Systems (HDS), IBM, NetApp, Red Hat, Scality, StorageCraft, SUSE, SwiftStack, and Western Digital (WD).

To make this list, the suppliers need to have an OBS platform that:

- **Conforms to IDC's taxonomy on OBS platforms.** According to IDC's taxonomy, software-based OBS platforms can run on any commodity x86 platforms and do not have any specific hardware customizations (like custom ASICs or SoCs) mated to the software stack, and they leverage an OBS data organization scheme.
- **Has been developed in-house or owned by way of an acquisition.** In other words, the supplier needs to be the intellectual property owner of that platform.
- **Is delivered as software, hardware (appliance or gateway), and/or as (private or public) cloud based.** Additional points were granted to solutions supporting data tiering to and from public clouds through a simple intuitive UI as well as those that supported search and data analytics reporting capabilities.
- **Is sold as licensed software directly to buyers or indirectly via OEM/channel partners and not just as a service.** Additional points were granted if the supplier had partnerships with as-a-service providers to deliver it as a cloud offering.
- **Was generally available (GA) as a current offering at the time IDC undertook this study in early 2018 with revenue of $20 million+.**

This study is designed to evaluate each supplier for its OBS offering as opposed to the breadth of products and services of the firm. In other words, it should be observed that this study evaluates each participating supplier as an entity within the OBS market.

In addition, some suppliers did not make the list because they did not meet one or more of the selection criteria. The supplier that gets an honorable mention in this study is DataDirect Networks. Other suppliers in adjunct markets such as data protection are also mentioned.

Please see the "Reading an IDC MarketScape Graph" section for more details on how market share was calculated for vendors with multiple product offerings and delivery modes.

ADVICE FOR TECHNOLOGY BUYERS

As end users move toward a hybrid cloud/multicloud storage strategy to support the needs of their ever-increasing data sets while supporting traditional and next-generation workloads, IDC urges you to think about the following when choosing OBS solutions:

- **Automated information life-cycle management (ILM):** Policy-driven ILM capabilities that enable management, repair, and deletion of data will help increase efficiency when managing petabyte- or exabyte-scale data sets.
- **Solutions portfolio:** A strong technology (hardware and software) partnership portfolio lends itself to a given OBS offering being able to support many use cases across verticals. An OBS supplier with several ISV and server-hardware partnerships provides end users the ease of procuring and deploying solutions in a quick and efficient manner.
- **Data access:** A flexible OBS solution will allow native data access capability for files and objects where data can be stored or accessed as file or objects or native public cloud format.
- **Data tiering:** Data can be tiered to multiple public cloud storage services while being managed through a single pane of glass giving customers choice and ease of management to use the right platforms for their workloads.
- **Ability to support production workloads:** OBS solutions are slowly but surely being deployed for production use cases. As customers look at OBS solutions, it is imperative that they consider solutions that can support the performance, bandwidth, scale, and other requirements of production workloads.

VENDOR SUMMARY PROFILE

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

NetApp

NetApp is positioned as a Leader in this IDC MarketScape for object-based storage.

NetApp's entry in the OBS space, with its offering — NetApp StorageGRID — is via acquisition. Since its introduction in 2014, NetApp has made strides in improving and expanding the capabilities of NetApp StorageGRID. Today, NetApp's OBS is a part of the company's Cloud Infrastructure business unit that includes its hyperconverged, converged, and Active IQ offerings. The product is available as either an appliance or a software-only offering and offers enterprise-grade attributes such as client connectivity via known protocols (NFS/CIFS, cloud protocols like S3 and Swift), massive scale, global namespace across sites, comprehensive and automated ILM/dynamic policy management, flexible data protection methods (replication and erasure coding), fault-tolerant architecture, security, audit capabilities, and integration with cloud and archiving solutions.

In recent times, NetApp StorageGRID has made significant improvements to its integrated ILM functionality by providing controlled data placement and accessing policies by region using location constraint and updates to GUI for ease of use. The product also supports a new public metrics API that enables users to extract value from metrics collected at the node and service level for better insights. StorageGRID Cloud Tiering and Cloud Mirroring support data tiering to S3-compatible public clouds. It supports metadata search, reporting, and visualization integration across on-premises and cloud

deployments via the Elastic Stack. Search is also enabled for policy management via the product's ILM framework. StorageGRID now provides cloud-architected infrastructure for financial and personal data retention compliance as one integrated resource across public and private clouds.

As NetApp StorageGRID evolves in its functionality, it has developed a portfolio of offerings to better address the needs of the market. For example, NetApp StorageGRID NAS Bridge, the company's self-built cloud gateway, supports NAS protocols where files can be ingested as SMB or NFS or as objects via S3, targeted primarily at ROBO and archive use cases. NetApp's Fabric Pool allows customers the choice of leveraging Data ONTAP's file capabilities while tiering data to NetApp StorageGRID. In addition, NetApp's AltaVault and StorageGRID provide optimized solutions with file capabilities for backup and archive applications.

Strengths

The strength of NetApp comes in the form of its core engineering strengths and its vast experience in the unstructured data space. NetApp has made strides in bringing to market new products across market segments and partnership with hyperscale cloud storage service providers Microsoft Azure and Amazon AWS. The company's strategic vision to recognize the needs of the market and the willingness to drive effective change within its organization to align to its vision are commendable. With these changes, NetApp StorageGRID is a part of a broader integrated portfolio offering that makes it an attractive proposition to customers.

Challenges

In recent times, NetApp has increased its focus on new high-growth segments of the market such as all-flash arrays and hyperconverged infrastructure (HCI). Although NetApp continues to invest in engineering efforts to harden its OBS product, the company will need to focus on highlighting the value and advantages of NetApp StorageGRID to gain mindshare and market share. Therefore, NetApp's StorageGRID's biggest challenge is that of brand recognition in the OBS space.

APPENDIX

Reading an IDC MarketScape Graph

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

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

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

The size of the individual vendor markers in the IDC MarketScape represents the market share of each individual vendor within the specific market segment being assessed.

In the case of a supplier with multiple products in the same market segment, IDC has worked with the supplier to select the product that most closely resembles the tactical strengths (capabilities) and strategic directions (strategies) of the supplier and the one that can be used as the lens through which the supplier's position in the market can be ascertained, provided the product meets the inclusion criteria for the IDC MarketScape. This can impact the size of the bubble as only the revenue for the evaluated product is included and not the supplier's overall revenue for that market segment.

Therefore, while certain suppliers are at an advantage given their size and broader portfolio offerings, IDC recognizes that smaller suppliers with a single product, and whose primary focus in the object-based storage market may be limited to specific verticals, also play an important role by bringing to market potentially disruptive technologies.

Note that certain suppliers (e.g., Scality and Red Hat) are pure-play software vendors, while the other suppliers sell a mix of hardware and software, mostly as hardware appliances. Pure-play software typically represents 25–50% of the total revenue, so associated server revenue is added to compare the size of the bubbles directly to the appliance vendors.

IDC MarketScape Methodology

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

Market Definition

According to *Worldwide File- and Object-Based Storage Forecast, 2016-2020* (IDC #US41685816, September 2016), object-based storage capacity is expected to grow at a CAGR of 30.7% for 2016-2020, reaching 293.7EB in 2020.

IDC classifies OBS platforms as part of the scale-out file- and OBS (FOBS) market segment. IDC uses the classification scheme to classify newer software-based file- and object-based storage platforms.

Scale-out FOBS refers to FOBS solutions that use a distributed data placement mechanism to span multiple independent server hosts or controllers while presenting a single data access namespace. Such architectures are also called shared nothing (or sharded data) architectures. Such architectures allow for flexible scalability in performance and capacity independent of each other using commodity components. Data sharding and distribution mechanisms (such as local and geographic replication and local and distributed erasure coding) account for one or more concurrent component failures. Scale-out FOBS solutions are made up of two variants: scale-out FBS solutions and scale-out OBS solutions. There are two principal differences between the two types: how data is organized, and how data is accessed.

Scale-out FBS solutions use distributed file systems with hierarchical structures to organize and store data. These structures are akin to mechanisms used by monolithic file systems, which in most cases follow a root directory (folder) and inverted tree structure. In contrast, scale-out OBS solutions use flat

structures to organize data. Such structures are higher-level structures in which data is often organized using an "account, container, object" approach wherein "objects" are analogous to "files" in FBS solutions. Accounts, containers, and objects are referenced by a metadata repository that stores and manages attributes of data stored in that structure. The level at which OBS solutions operate varies from platform to platform. Many OBS solutions operate on a per-object level (i.e., allow each object to be treated independently as far as policy management is concerned), whereas others operate at a container or account level (i.e., only allow policies to be applied at a container or account level). Several OBS solutions also leverage NoSQL databases as metadata repositories and persistent data stores (instead of storing chunks in the file systems).

Because of the need to manage objects with a comprehensive set of attributes, most OBS solutions use a different set of data interfaces than their FBS counterparts that mostly leverage NFS, SMB (CIFS), or FTP protocols. It is common for many OBS solutions to support HTTP/REST, CDMI, Amazon S3, and other object-specific interfaces.

IDC's File- and Object-Based Storage User Survey Findings

In 2017, IDC conducted a survey of 450 concurrent users of file and object-based storage in North America. The study, called IDC's *File- and Object-Based Storage Survey Findings,* reported:

- 52.9% of 450 North America–based respondents started using OBS in the past two years.
- 41% of the respondents also indicated that they are currently using either an OBS appliance or a commercial software-only offering, while 42% are using public cloud or managed private cloud services.
- Over 26% of 254 respondents plan to integrate flash into their OBS environment, while 56.4% have already done so. The primary reasons to integrate flash are higher storage density, bandwidth (to move large data sets faster), and performance (throughput). Flash is integrated into OBS to support data analytics of unstructured data, content applications, business applications, and others.
- 48% of 281 respondents indicate that they use OBS in the public cloud to support a new initiative, while 42.7% indicate (lower) cost as a secondary reason to adopt OBS in the public cloud. Better data security and data protection are other top reasons to adopt public cloud.
- 422 respondents deployed OBS either as traditional on-/off-premises or as private cloud storage. 37.7% of users deploying traditional on-/off-premises OBS and 34.6% of users deploying private cloud indicate the reason for this choice of platform is because of sensitivity of data (cannot be pushed into public cloud).

The research also indicates that the actual benefits of deploying OBS in terms of avoiding vendor lock-in and flexibility of platform (private/public cloud, appliance, software only, open source software, etc.) are greater than perceived benefits.

Backup, DR, Archive, and Active Archive are still dominant use cases for OBS. Newer use case such as data analytics for unstructured data, media streaming, and web serving are deployed on off-premises OBS and are expected to be revenue-generating workloads in the future.

LEARN MORE

Related Research

- *IDC's Worldwide SBS, SDS, and FOBS Storage Solutions Taxonomy, 2018* (IDC #US43579118, March 2018)
- *Business and Content Applications Will Dominate Evolving Production SDS Deployments* (IDC #US43249017, December 2017)
- *Veritas Vision 2017 - Renewed Focus on Multicloud Data Management* (IDC #IcUS43114717, September 2017)
- *Worldwide Software-Defined Storage Forecast, 2017-2021: SDS Market Growth Significantly Outpaces Enterprise Storage Growth, Led by HCI* (IDC #US43062517, September 2017)
- *Worldwide File- and Object-Based Storage Forecast, 2017-2021* (IDC #US42280717, September 2017)
- *Hedvig Brings Hyperscale Agility and Economics to Enterprises* (IDC #US42983117, August 2017)
- *Successful SDS Deployments Demand Requisite Expertise on the Part of Customers* (IDC #US42919517, August 2017)
- *IDC's Worldwide Storage Software Taxonomy, 2017* (IDC #US42834017, August 2017)
- *Veritas HyperScale for OpenStack: Consistent Performance and Data Protection for OpenStack-Based Cloud* (IDC #US42580517, June 2017)

Synopsis

This IDC study represents a vendor assessment model called the IDC MarketScape. This study is a quantitative and qualitative assessment of the characteristics that assess a vendor's current and future success in the said market or market segment and provide a measure of the vendor's ascendancy to become a leader or maintain leadership. IDC MarketScape assessments are particularly helpful in emerging markets that are often fragmented, have several players, and lack clear leaders.

The (scale-out) OBS market subsegment, which is part of the file and OBS market, is an example of an emerging market. In this IDC MarketScape, IDC attempts to assess the capabilities and strategies of key vendors of OBS solutions. IDC expects that market forces such as fierce competition and buyer demand will accelerate the metamorphosis of this market into a mature market with only a few dominant vendors. Open source-based stacks will create an additional dimension of complexity and challenges. In all likelihood, the only survivors in this market may be vendors with robust partner ecosystems and/or vendors with commercial variants of open source platforms.

"A new digitized world demands an infrastructure that is extremely scalable and flexible in terms of delivery models and also payment options, with strong vendor strategic and research vision from the vendor side along with unprecedented economies of scale," said Amita Potnis, research manager in IDC's Storage team. "OBS platforms hold the promise and the potential to support end users along this path of digitization. In this competitive market, vendors offering OBS platforms with the most compelling value proposition via a long-term strategy, research and development plan, and flexible delivery models will survive."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com