

## IDC MARKETSCAPE

# IDC MarketScape: Worldwide Object-Based Storage 2014 Vendor Assessment

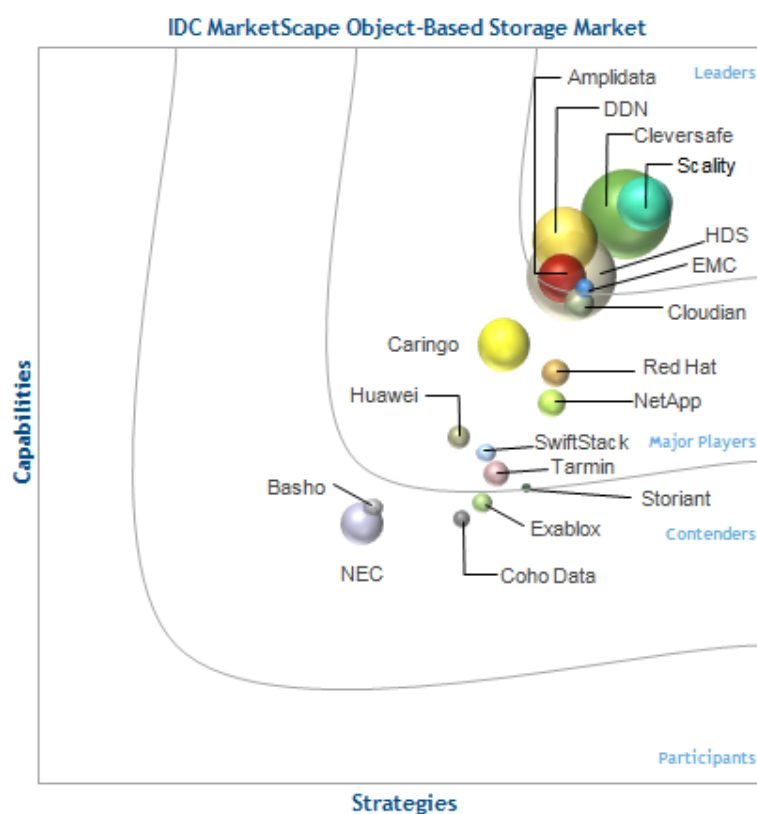
Ashish Nadkarni

THIS IDC MARKETSCAPE EXCERPT FEATURES: CLEVERSAFE

## IDC MARKETSCAPE FIGURE

FIGURE 1

### IDC MarketScape Worldwide Object-Based Storage Vendor Assessment



Note: Suppliers include ISVs and appliance providers. Please see the Appendix for detailed methodology, market definition, and scoring criteria. In the case of suppliers with multiple products (e.g. EMC), only the revenue for the selected product is referenced. This impacts the size of the bubble.

Source: IDC, 2014

## IN THIS EXCERPT

---

The content for this excerpt was taken directly from IDC MarketScape: Worldwide Object-Based Storage 2014 Vendor Assessment (Doc #253055). All or parts of the following sections are included in this excerpt: IDC Opinion, IDC MarketScape Vendor Inclusion Criteria, Essential Guidance, Vendor Summary Profile, Appendix and Learn More. Also included is Figure 1.

## IDC OPINION

---

Object-based storage (referred to as OBS throughout this document) platforms continue to perpetuate cloud and enterprise IT infrastructure. As businesses move toward petabyte-scale data storage, OBS solutions are turning out to be the right choice for balancing scale, complexity, and costs. By way of their core design principles, OBS platforms deliver unprecedented scale at reduced complexity and reduced costs over the long term. Early OBS platforms suffered from "necessity crisis," were too cumbersome to deploy and, in some cases, caused a platform lock-in because of their proprietary access mechanisms. In spite of their from-the-ground-up design, a departure from how traditional SAN and NAS arrays are deployed and, more importantly, a lack of standard interfaces made it difficult for IT organizations to deploy OBS solutions in the infrastructure. Thanks to Amazon S3 and OpenStack Swift becoming de facto access interfaces, this situation has changed. This has in turn spurred further investments in making OBS solutions more ubiquitous. In 2014, several new suppliers entered the market, while others announced partnerships and expanded their ecosystem, and notably, many of them came together to form an alliance (known as the Object Storage Alliance) to perpetuate OBS further into the enterprise. An important trend that is catching on slowly is to deemphasize the "OBS" component of the solution and instead focus on the use case or the problems solved by the solution – for vertical-specific situations or for more generic Big Data, social, cloud, and mobile workloads. IDC estimates that in 2014, OBS solutions accounted for nearly 46% share of the file-and-OBS (FOBS) market in revenue and is forecast to be a \$28.3 billion market in 2018. However, the total revenue from commercial OBS products as a percentage of the overall OBS market remained quite small in 2014, and even with a forecast CAGR of over 27%, it will constitute a relatively small portion in 2018. IDC estimates that in 2014 the "all-in" size of the commercial object storage systems market in terms of revenue – which includes hardware and software from OBS suppliers and their OEM partners – was around \$775 million. In this IDC MarketScape, IDC assesses the present commercial OBS supplier (suppliers that deliver software-based OBS solutions as software or appliances much like other storage platforms) landscape. As this IDC MarketScape indicates, in 2014, the OBS market became more fragmented (IDC now counts more than 25 OBS suppliers, of which 18 are featured in this IDC MarketScape as "owner of IP" suppliers). IDC analyzed the capabilities and business strategies of OBS suppliers that it considers to be representative of the market. Key findings include:

- The market has more or less adopted Amazon S3 and OpenStack Swift as the RESTful access mechanisms for object storage. Several suppliers are aligning themselves to OpenStack in an attempt to make their object storage compatible with Swift, Cinder, Manila, Glance, and Nova. They are also trying to preemptively ward off a threat from Ceph and Swift (the OpenStack OBS implementation). This has also led to a new breed of open source-based commercial variants of OBS platforms like Swift and Ceph (some of which are profiled in this IDC MarketScape).
- Leading suppliers (especially incumbent storage suppliers) are focusing on a broader portfolio and a bigger sales force to go after more buyers. Furthermore, many such suppliers have

signaled a portfolio approach wherein object capabilities would be provided using multiple platforms – both legacy and next generation.

- Smaller suppliers (especially the newer entrants) that have more innovative solutions are getting additional funding to go after a diverse buyer base. Such suppliers therefore focus on specific verticals and/or use cases such as archiving, telecom providers, and managed service providers as they build out their business. They are also extending their reach using OEM partnerships and reseller relationships.

## IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

---

This IDC study assesses the capabilities and business strategies of leading suppliers in the (scale-out) OBS market segment – which is part of the overall file-and-OBS 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 the 18 most notable players in the commercial OBS market, with broader portfolios and global scale. The suppliers enlisted in this study are (in alphabetical order): Amplidata, Basho Technologies, Caringo, Cleversafe, Cloudian, Coho Data, DataDirect Networks (DDN), EMC, Exablox, Hitachi Data Systems (HDS), Huawei, NEC, NetApp, Red Hat (Inktank), Scality, SwiftStack, Storiant, and Tarmin.

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 now owns the platform by way of an acquisition.** In other words, the supplier needs to be the intellectual property (IP) 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 software-based platforms that can be installed on any commodity-based hardware.
- **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 2014.**

This study is designed to evaluate each supplier for its (scale-out) 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 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 direction (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 MarketScope. 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.

While certain suppliers therefore 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 scale-out 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, Red Hat, and Basho) 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 must be added to compare the size of the bubbles directly to the appliance vendors.

In addition, some suppliers did not make the list because they did not meet one or more of the selection criteria. The suppliers that are mentioned in the study are Quantum, HGST, Avnet, and Nexenta. Other suppliers in adjunct markets such as data protection are also mentioned.

## ESSENTIAL BUYER GUIDANCE

---

All companies, small and large, grapple with data growth. As businesses become data driven to survive in the new economy, they will seek more data sources, collect more data, and look to analyze and store this data in a decentralized manner. In many cases, they will look to perform real-time analytics on this data as it is generated and where it gets captured. Many others will seek to create on-demand opex-driven cloud environments for internal and external consumption. Nontraditional use cases, especially for highly scalable and decentralized semistructured (machine generated) and unstructured data storage, will require nontraditional storage solutions like OBS. OBS is an innovative approach to storage, and its procurement needs to be preceded by careful planning. Unlike traditional SAN or NAS arrays, OBS solutions are not built the same way nor are they built to suit all use cases and workloads equally. Buyers should therefore look for the following key characteristics when evaluating OBS solutions:

- **Platform scalability:** Scalability is not just from a hardware perspective but also from throughput, file size, and file volume perspectives. A solution appropriate for a given environment will allow each dimension to scale independently.
- **Data management:** Data layout and organization is an important piece as it may have performance, efficiency, and availability implications. Over time, as data grows, organizations will face the need to mine existing data for patterns that may build new business cases around new findings. A solution that supports advanced metadata, indexing, and analytics will be a key component of the infrastructure.
- **Storage efficiency:** The larger the data set and bigger the storage system, the greater the need of data management and reduction techniques (data deduplication, compression, thin provisioning, etc.). Data optimization technologies (automated data tiering) will also be essential. A solution appropriate for a given environment will allow many, if not all, of the previously mentioned features to be implemented and recalibrated without major disruptions.
- **Data resiliency:** Resiliency capabilities (like replication and erasure coding) and the granularity with which such capabilities can be applied (i.e., whether policies can be applied at an account, container, or object level) will be important considerations. Data resiliency should

also be weighed against the platforms' CAP theorem profile (see the Appendix for more details).

- **Workload adjacency:** Several OBS platforms offer or are considering offering nonstorage workloads to run natively on the OBS platform. Since most OBS platforms are node based and use x86 platforms, they offer excellent workload adjacency for distributed and localizable workloads like Map/Reduce and hypervisors. This is an important consideration for decentralized storage and in situations where the data has a short shelf life.

In addition to the platform characteristics, buyers should look for the following supplier attributes:

- **Supplier's commitment to the platform now and for the future:** Strong road map, customer support and service, and overall track record on incorporating new features into the platform are some of the attributes buyers should look for in a supplier.
- **Partner ecosystem for applications and on-ramping:** The more comprehensive the ecosystem, the better placed the supplier in offering an end-to-end workload-optimized or use case-focused solution.

This study, while an independent evaluation from a market research firm, should not be considered as a "final judgment" on suppliers when considering their respective offerings for a particular project. The specific objectives and requirements of any end-user company will play a significant role in determining the suppliers that can be considered as candidates for an engagement.

## VENDOR SUMMARY PROFILES

---

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

### Cleversafe

Cleversafe was founded in 2004, with private investments, and is headquartered in Chicago, Illinois. The company is aimed at building solutions that have inherent features such as new levels of scalability, security of data, storage efficiency, availability, and reliability, with the option to deploy across geographies or within a single datacenter. With a proven ability to scale to petabytes and beyond, Cleversafe's solutions promise extreme scalability and simplicity, combined with data protection and security, at significantly lower cost than solutions relying on storing more than one copy of the data to meet reliability requirements. Features such as these have won Cleversafe contracts with government agencies for defense and security purposes.

Cleversafe's flagship product is its Dispersed Storage Network (dsNet) solution, which slices and disperses unstructured data across multiple nodes or geographical regions using patented Information Dispersal Algorithms. The dsNet platform can be deployed as a software solution on existing infrastructure, as a hardware appliance sold by Cleversafe, or on qualified third-party hardware. The dsNet is made up of three components: dsNet Manager, Accesser, and Slicestor. The dsNet Manager enables operations and maintenance for the object repository, whereas Slicestors are the storage nodes. Accessers provide a REST interface to applications and have the intelligence to encrypt and

encode the data on writes and reverse the process on reads. This erasure code-based encoding protects the data against drive, server, rack, and site failures by transforming it into a number of slices such that the original data can be read using a subset of those slices. This approach is more efficient than data protection schemes that combine RAID and replication. One of Cleversafe's major selling points is the inherent security and reliability of its platform – it includes SecureSlice, a keyless encryption technology. This feature allows objects to be encrypted before being sliced and dispersed onto the Slicestor(s). Data can only be accessed when the required threshold of slices have been retrieved. In addition, no one Slicestor has all of the object's slices, thus adding another security layer. To maintain data integrity at scale and over time, a multilevel integrity checking scheme and background data repair process are built into the product. Users can configure fault tolerance levels depending upon the number of sites, number of nodes, and size of their configuration. Cleversafe uses both direct and indirect sales channels focusing on specific regions and verticals to drive business. Cleversafe has also made several technological partnerships through its Solution Network Program. Some of these partnerships include gateway providers such as Panzura, Nasuni, Avere, Riverbed, and CTERA whose products expose CIFS and/or NFS to client application and technologies such as Symantec Netbackup, CommVault Simpana, and QStar's Archive Manager – backup and archiving software that enables data protection and long-term retention of objects on Cleversafe's storage.

IDC has placed Cleversafe in the Leaders category in this IDC MarketScape. This position reflects Cleversafe's current capabilities and a strategy to continue leading this market. Cleversafe has done well in focusing on what it knows best – massive at-scale deployments best served by OBS platforms.

## **Strengths**

Cleversafe has the advantage of being an early entrant with next-gen technology in the OBS market segment. This has allowed it to gather several years of experience early on. Some of Cleversafe's earliest customers were government agencies with rigid requirements. Working with them, Cleversafe gained the experience necessary to enter into the enterprise while also taking on some of the hyperscale and Web 2.0 firms as its clients. Cleversafe's strengths lie in seeing the need for OBS platforms early on and successfully maneuvering past the company's early challenges to successfully make headway into newer verticals such as media and entertainment, healthcare, and managed service providers.

## **Challenges**

Cleversafe is on an aggressive growth path, with projections to more than double its revenue and customer base in 2014 for the third year in a row. It needs to continue to make the most out of its position in the market to lead the way when it comes to next-generation OBS platforms for private and public cloud deployments. By taking a more active role in alliances like the Object Storage Alliance, it can foster a community of OBS platform vendors and their ecosystem partners. The company is continuing to expand geographically beyond its focus on North America and select portions of EMEA and APAC and needs to use its HP relationship to accelerate expansion. By leveraging its leadership position and partner opportunities, Cleversafe will no doubt emerge as a dominant force as and when this market segment matures.

## 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 18 months.

The size of the individual vendor markers (i.e., bubbles) in the IDC MarketScape represents the market share of the evaluated product of each individual vendor within the specific market segment being assessed.

From a product currency and future perspective, IDC examined the core functionalities offered by the product. These included data organization, encoding/decoding mechanisms, data resiliency mechanisms (like replication and erasure coding), support for flash, geo-dispersal/redundancy capabilities, local data persistence (DB/FS/Raw), metadata format, multitenancy, native data management capabilities, support for native scale-out file system, parallelism, tiering, topology, and interfaces supported. In the case of interfaces supported, IDC asked vendors to list which interfaces their product supports or will support. Examples of such interfaces are CDMI, CIFS, HTTP/Rest, iSCSI, NFS, OpenStack Swift and Cinder, OST, POSIX, and of course Amazon S3.

Some of the use cases (and application integrations) cited by vendors for their platforms were email, Native Hadoop, media access manager (MAM), iRODS, sync-n-share apps, content repositories, archiving, regulatory and compliance archive, backup reduction/optimization, storage and server consolidation, secure cloud storage, primary storage and/or primary storage optimization, governance, risk and compliance applications, and enterprise search. Similarly, vendors noted the following as the verticals with the most traction for their products: healthcare, energy and gas, media and entertainment, Internet content delivery, security/antivirus, enterprise information archive, ECM or ERM, eDiscovery, ERP, education, public sector/government, life sciences, managed service and telecom service providers, and BPO/back office.

### IDC MarketScape Methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and selected 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 a review board of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions in 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.

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.

## Market Definition

IDC classifies OBS platforms as part of the scale-out file-and-OBS market segment. IDC uses the classification scheme illustrated in Figure 2 to classify newer software-based file-and-object platforms.

**FIGURE 2**

### IDC Software-Based Storage Classification Scheme



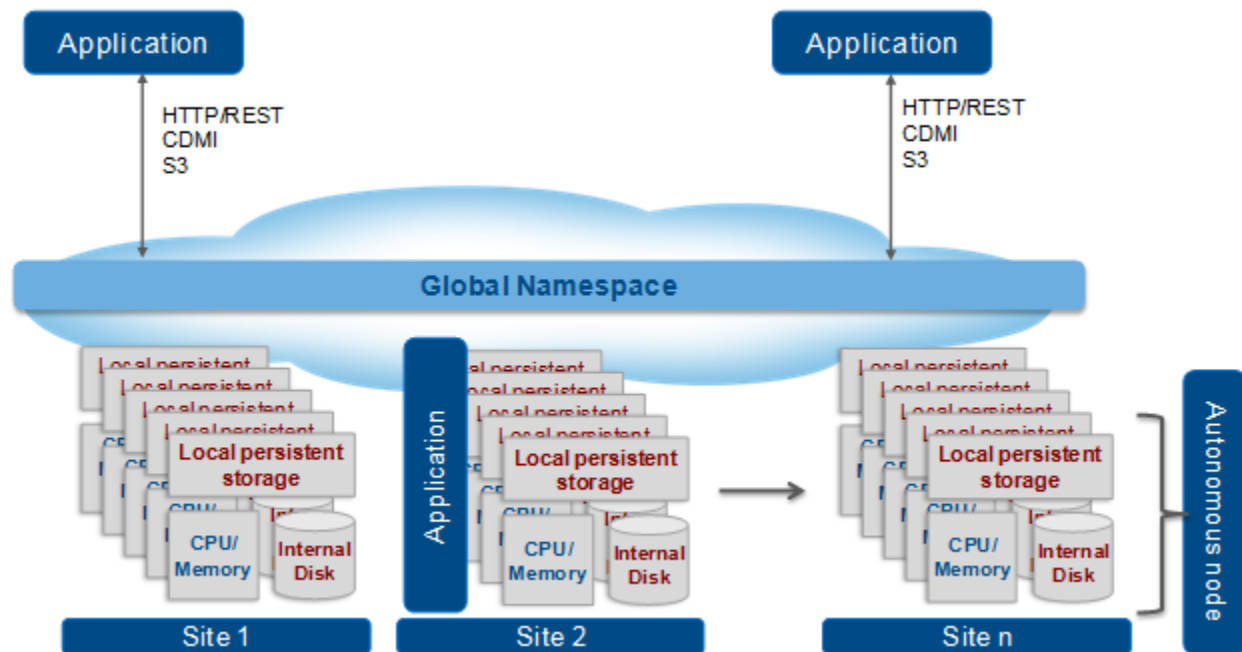
Source: IDC, 2014

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 and are illustrated in Figure 3. 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.



**FIGURE 3**

### Shared Nothing (Shared Data) Architectures



Source: IDC, 2014

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 follows 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.

### Related Research

- *Worldwide File- and Object-Based Storage 2014-2018 Forecast* (IDC #251626, October 2014)
- *IDC's Worldwide File- and Object-Based Storage Taxonomy, 2014* (IDC #245940, January 2014)
- *Worldwide File- and Object-Based Storage 2013-2017 Forecast* (IDC #242287, July 2013)
- *Ceph: An Emerging Converged Platform for the Enterprise* (IDC #242140, July 2013)
- *Red Hat Summit Message: All-In on OpenStack* (IDC #1cUS24189013, June 2013)
- *Worldwide Enterprise Storage Systems 2013-2017 Forecast: Customer Landscape Is Changing, Defining Demand for New Solutions* (IDC #241033, May 2013)
- *OpenStack: Why It Matters and to Whom* (IDC #240517, April 2013)
- *IDC's Worldwide Software-Based (Software-Defined) Storage Taxonomy, 2013* (IDC #240500, April 2013)
- *Shared Nothing Architectures – A Blueprint for Software-Based Scale-Out Solutions* (IDC #239526, February 2013)
- *IDC's Worldwide File- and Object-Based Storage Taxonomy, 2013* (IDC #239143, January 2013)
- *Worldwide Storage 2013 Top 10 Predictions: Market Transformation from Convergence, Cloud, and Scale-Out Strategies* (IDC #238996, January 2013)
- *Scality: Defining Unified Software-Based Scale-Out* (IDC #238458, December 2012)
- *Worldwide Enterprise Storage Systems 2012-2016 Forecast Update* (IDC #237886, November 2012)
- *Red Hat as a Storage Solutions Provider* (IDC #237976, November 2012)
- *How Distributed File Systems Are Rewriting the Future of the Storage Ecosystem* (IDC #236517, August 2012)
- *Worldwide File-Based Storage 2012-2016 Forecast: Solutions for Content Delivery, Virtualization, Archiving, and Big Data Continue to Expand* (IDC #235910, July 2012)
- *The Future for Namespaces in File-Based Storage* (IDC #236010, July 2012)
- *Worldwide Storage and Virtualized x86 Environments 2012-2016 Forecast* (IDC #235868, July 2012)
- *Worldwide Enterprise Storage Systems 2012-2016 Forecast* (IDC #234990, May 2012)
- *Worldwide Storage and Virtualized Environments 2011-2015 Forecast* (IDC #231080, December 2011)
- *Worldwide Enterprise Storage Systems 2011-2015 Forecast Update* (IDC #228255, May 2011)

- *Worldwide Storage and Virtualized x86 Environments 2010-2014 Forecast* (IDC #224344, August 2010)

## 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 their ascendancy to become a leader or maintain a 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 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.

"As businesses embark on a transformation to become data-driven entities, they will demand a data infrastructure that supports extreme scalability and flexible acquisition patterns and offer unprecedented economies of scale," said Ashish Nadkarni, research director in the IDC Storage Team. "OBS platforms hold the promise and the potential to assist buyers along this data-driven journey. Vendors offering OBS platforms with the most compelling value proposition will survive. Others will perish. This is a race to the finish."

For more details and clarification, please read IDC's clarification on this IDC MarketScape at [idc-community.com/groups/it\\_agenda/storageanddatamanagement/2014\\_marketscape\\_on\\_object\\_based\\_storage\\_platforms\\_clarification\\_on](http://idc-community.com/groups/it_agenda/storageanddatamanagement/2014_marketscape_on_object_based_storage_platforms_clarification_on).

## 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-insights-community.com](http://idc-insights-community.com)  
[www.idc.com](http://www.idc.com)

---

### Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit [www.idc.com](http://www.idc.com) to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit [www.idc.com/offices](http://www.idc.com/offices). Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or [sales@idc.com](mailto:sales@idc.com) for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2014 IDC. Reproduction is forbidden unless authorized. All rights reserved.

