



IDC MarketScape

# IDC MarketScape Excerpt: Worldwide Object-Based Storage 2013 Vendor Assessment

Ashish Nadkarni

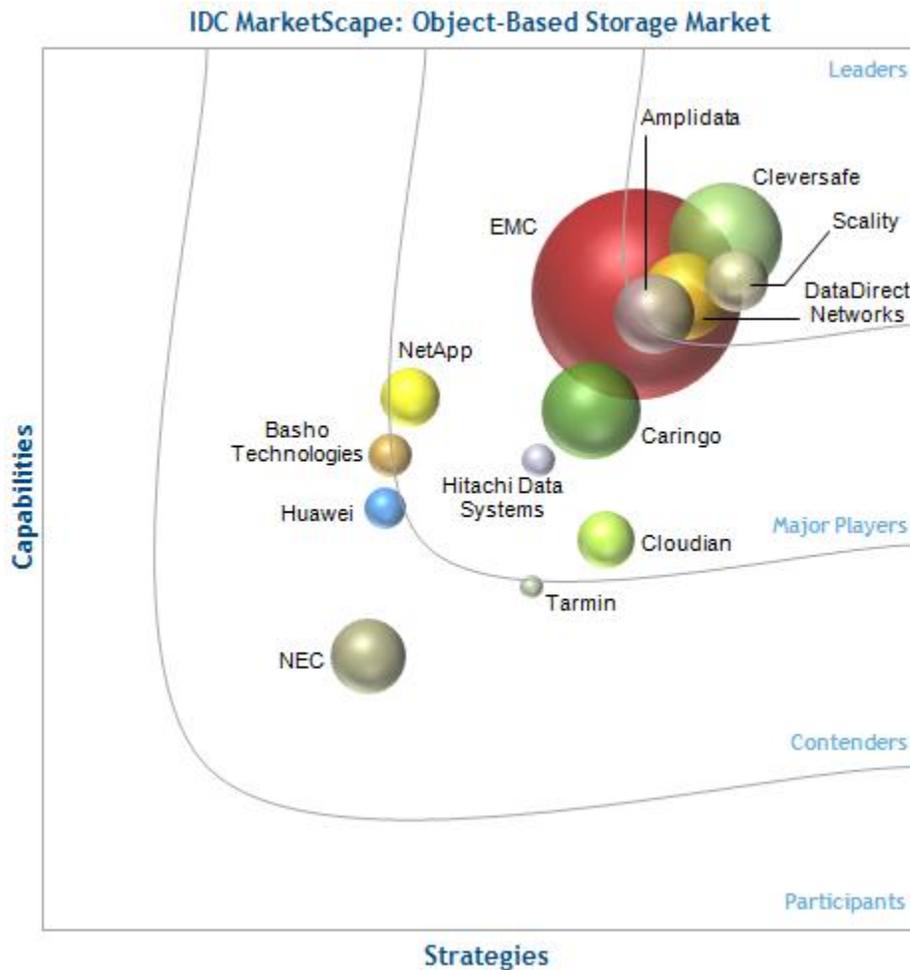
Amita Potnis

**THIS IDC MARKETSCAPE EXCERPT FEATURES: EMC**

## IDC MARKETSCAPE FIGURE

**FIGURE 1**

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



Source: IDC, 2013

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

## IN THIS EXCERPT

---

The content for this paper is excerpted from the IDC MarketScape: Worldwide Object-Based Storage 2013 Vendor Assessment, (Doc # 244081). 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. Figure 1, Figure 2, Figure 3 and Figure 4 are also included.

## IDC OPINION

---

Object-based storage (referred to as OBS throughout this document) platforms offer a promising future for organizations trying to balance scale, complexity, and costs. OBS platforms by way of their core design principles 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. This situation is however fast changing. In 2013, several new and existing vendors of OBS platforms have announced partner programs, joined alliances for standards, and made the vendors' offerings more purpose built – all in an attempt to prove that OBS platforms do indeed solve the problems that they were designed to solve, making them "must-have" for several use cases, workloads, and verticals. They have been helped in large part by open source communities that have developed OBS stacks for DIY deployments. OpenStack Swift in particular has gained a lot of traction both in the enterprise and in the service provider space. All these trends support IDC's assertion that 2013 marked a turning point for the OBS market. IDC estimates that in 2013, OBS solutions accounted for nearly 37% share of the file-and-OBS (FOBS) market in revenue and is forecast to be a \$21.7 billion market in 2017. However, the total revenue from commercial OBS products as a percentage of the overall OBS market is quite small in 2013, and even with a forecast CAGR of over 27%, it will constitute a relatively small portion in 2017. In this IDC MarketScape, IDC assesses the present commercial OBS vendor (vendors that deliver software-based OBS solutions as software or appliances much like other storage platforms) landscape. IDC analyzes the capabilities and business strategies of 13 OBS vendors that it considers to be representative of the market. Key findings include:

- Leading vendors (especially incumbent storage vendors) have a broader portfolio and a bigger sales force to go after more buyers. This in some cases leads to a bigger market share by revenue but does not necessarily imply a strong OBS platform in terms of capabilities. Furthermore, many of such vendors have signaled a portfolio approach wherein object capabilities would be provided using multiple platforms – both legacy and next generation.
- Smaller vendors (especially the newer entrants) have more innovative solutions but lack the financial muscle to go after a diverse buyer base. Many such vendors 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 also extend their reach using OEM partnerships and reseller relationships.

- All vendors acknowledge the need to make their solutions appeal to a wider buyer base and are aggressively moving their product capabilities in that direction. However, the lack of standards means that feature-to-feature comparison will continue to challenge buyers and make it difficult for them to select a platform. In many cases, buyers resort to the vendors' overall capabilities to aid their decision process.

## IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

---

This IDC study assesses the capabilities and business strategies of leading vendors 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 vendor to be successful in delivering an OBS solution in the market. This study includes analysis of the 13 most notable players in the commercial OBS market, with broader portfolios and global scale. The vendors enlisted in this study are (in alphabetical order): Amplidata, Basha Technologies, Caringo, Cleversafe, Cloudian, DataDirect Networks (DDN), EMC, Hitachi Data Systems (HDS), Huawei, NEC, NetApp, Scality, and Tarmin. To make this list, the vendors 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 vendor 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 vendor 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 2013.**

This study is designed to evaluate each vendor 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 vendor as an entity within the OBS market. In the event that the vendor has more than one OBS offering, IDC has selected the platform that most closely resembles those from other vendors included in this IDC MarketScope. So while certain vendors therefore are at an advantage given their size and broader portfolio offerings, IDC recognizes that smaller vendors, whose primary focus in the scale-out file-based storage market may be limited to specific verticals, also play an important role by bringing to market potentially disruptive technologies.

Additionally, some vendors did not make the list because they did not meet one or more of the selection criteria. The vendors that are mentioned in the study are Exablob, Inktank, Quantum, and SwiftStack. Other vendors 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. Non-traditional use cases, especially for highly scalable and decentralized semi-structured (machine generated) and unstructured data storage, will require non-traditional 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 section for more details.).
- **Workload adjacency:** Several OBS platforms offer or are considering offering non-storage 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 short shelf life.

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

- **Vendor'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 vendor.
- **Partner ecosystem for applications and on-ramping:** The more comprehensive the ecosystem, the better placed the vendor 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 vendors 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 vendors that can be considered as candidates for an engagement.

## VENDOR SUMMARY PROFILE

---

The section that follows briefly explains IDC's key observations resulting in one vendor's position in the IDC MarketScape. While every vendor is evaluated against each of the criterion outlined in the Appendix section, the description here provides a summary of one vendor's strengths and opportunities.

### EMC

EMC is one of dominant vendors in the storage industry, with a diverse portfolio that spans multiple market segments and encompasses entry, midrange, and high-end storage solutions across a wide variety of use cases and verticals. EMC entered the OBS market segment very early on with the Centera. Content-addressable storage, as it was known back then, could almost be considered to be the first-generation OBS that leveraged a node based, shared nothing architecture. In spite of its quirks and limitations, EMC quickly became a monopoly in this market segment. Subsequently, EMC introduced the Atmos platform, which was designed to be open, REST-based cloud storage, whereas the Centera was positioned as on-premises-based archive storage. Most recently, EMC also introduced ViPR, which is designed to provide virtualized object-based services leveraging any file block or OBS array and, eventually, commodity hardware.

Atmos is used to store, archive, and access unstructured data at scale, overcoming some of the limitations of the Centera. It is designed to enable businesses to take advantage of a private, public, or hybrid cloud storage environment. Atmos is delivered as a hardware appliance or software, and although it doesn't target specific verticals, it has experienced particular success with service providers, healthcare, oil and gas, financial services, and so forth. Atmos was organically developed by EMC and operates as a single entity. Atmos enables customizable metadata for data placement, protection, and life-cycle policies. Data protection is provided in two ways through erasure coding and also replication. The Atmos Cloud Delivery platform is an add-on software program that allows service providers or enterprises to deliver and manage storage as a service from their Atmos Cloud. Atmos can be accessed via a RESTful, CAS, or S3 interface and also exposes CIFS and/or NFS. EMC leverages a slew of channel partners that focus on various use cases that include service providers. Atmos hardware appliance is offered in three models – WS2-120 that scales to 60TB, WS240 supporting 720TB/rack, and finally WS-480 that scales to 1,440TB/rack. The most successful and widely known deployment of Atmos has been at AT&T. The introduction of Atmos Virtual Edition EMC can extend cloud capabilities using its other products. Atmos Virtual Edition can extend existing or new EMC storage or any VMware vSphere-certified storage like the Centera; EMC has developed an impressive partner ecosystem for Atmos. It includes companies like Aspera, a high-performance file-transfer solution; Cloudera for Hadoop environments; and Panzura for NFS/SMB access to the cloud.

IDC has placed EMC in the Leaders category in this IDC MarketScape. This position reflects EMC's commitment to this market segment, having established it early on with the Centera, then with Atmos, and in the future with ViPR.

## Strengths

EMC is the only incumbent storage vendor to be featured in the Leaders category. It is also the only vendor with multiple OBS solutions – Centera, Atmos, and now ViPR. Furthermore, EMC also offers object interfaces on the Isilon platform. EMC has built an impressive partner ecosystem for Atmos – which it is extending to the ViPR platform. EMC's unparalleled sales, marketing, and services organizations, which are no doubt helping execute on EMC's vision for the OBS segment, are the company's other strengths. EMC's commitment to this market segment is evident from the success Atmos has seen in private and public cloud deployments. The revenue for Atmos is growing at a rapid pace, and as is apparent (refer back to Figure 1), it's already much higher than all other competitors.

## Opportunities

One of the reasons Atmos has been successful is because it had little product overlap with another EMC solution. With ViPR – EMC's next-generation storage virtualization solution for block, file, and OBS – that situation has changed. EMC now finds itself in the position of having to justify multiple products in a segment in which the revenue from its OBS platforms is relatively small compared with its overall revenue from other storage platforms. As a strategy, EMC is positioning ViPR as its future object platform – Atmos and the Centera will continue to be portfolio products (with road maps that span the next couple of years). For EMC, ViPR will form the virtualization layer that provides universal object functionality. EMC should take this opportunity to educate the market on how these products will coexist and, as appropriate, which verticals and use cases each of them is best suited for. EMC should seek to extend its success with the Centera to Atmos and subsequently ViPR. Similar to how it established the CAS market and pretty much dominated it all the way, it can use Atmos and ViPR to dominate the OBS market segment. It should reinforce its commitment to maintaining the portability and openness of the platform by supporting all standards-based interfaces. Current and future buyers will no doubt appreciate and reward EMC in return.

## APPENDIX

---

### Reading an IDC MarketScape Graph

The IDC MarketScape vendor assessment for OBS market represents IDC's assessment on vendors that are well positioned today via the current capabilities of their OBS solutions and likewise vendors that are best positioned to gain market share over the next few years. Positioning in the upper right of the grid indicates that vendors are well positioned to gain market share. For the purposes of discussion, IDC divided potential key strategy measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis (or capabilities axis) reflects the alignment of the capabilities of the vendor and the vendor's corresponding OBS solution to current market demand. Via the capabilities category, IDC examines how well a vendor has aligned its solution to market demand. The further along a vendor is on this axis, the more successful the vendor is on this alignment. IDC considers the capabilities score to be a measure of the vendor's success in tactical execution for the OBS platform.

Positioning on the x-axis (or strategies axis) reflects the alignment of the future strategies of the vendor and the vendor's corresponding OBS solution to forecast market demand (three to five years). Via the strategies category, IDC examines how well a vendor plans to align its solution to its anticipated market demand. The further along a vendor is on this axis, the more successful the vendor is on this alignment. IDC considers the strategies score to be a measure of the vendor's vision and strategy to be a Major Player and/or a Leader in the OBS market segment over the next 18 months. The strategies category focuses on key areas such as strategic decisions and underlying assumptions about offerings, customer segments, business, platform road maps, and future go-to-market plans over an 18-month time frame.

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, multi-tenancy, 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), 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.

## 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

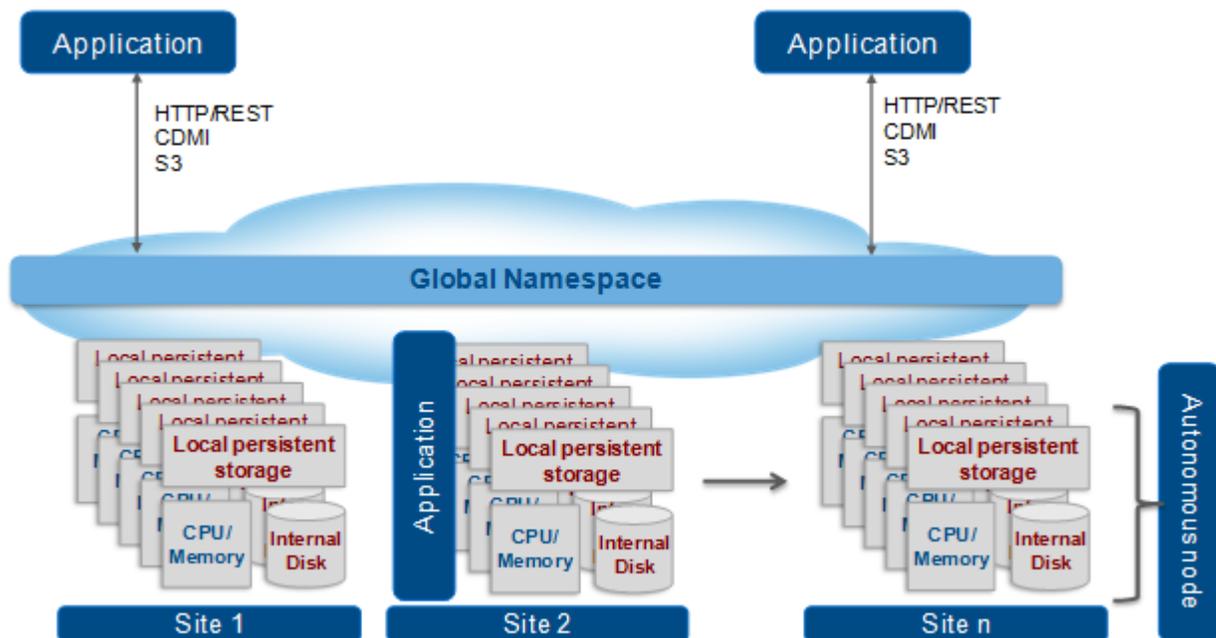


Source: IDC, 2013

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 (Sharded Data) Architectures



Source: IDC, 2013

Scale-out FBS solutions use distributed file systems with hierarchical structures to organize and store data. These structures are akin to the same 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.

## Use of CAP Theorem for Profiling OBS Platforms

The CAP theorem is often used to describe data consistency and integrity policies of OBS platforms. The CAP theorem states that it is impossible for a distributed computer system to simultaneously provide all three of the following guarantees:

- **Consistency:** The ability of all nodes to see the same data at the same time, regardless of updates or deletes
- **Availability:** A guarantee that every request receives a response about whether the request was successful or not (i.e., failed)
- **Partition tolerance:** The ability of the system to continue to operate despite failure of part of the system or lack of communication between nodes (also known as split brain)

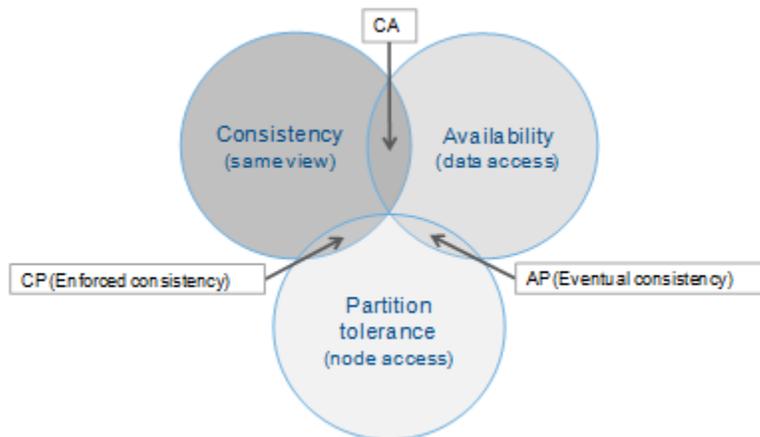
NoSQL and RDBMS/SQL databases often get measured using this theorem. Because of the architectural similarities, many now apply the CAP theorem to OBS platforms as well. For example, in OpenStack Swift, availability and partition tolerance are guaranteed but consistency is not. That means in Swift, one will always get the data; it will be dispersed on many places to make it highly available, but one could get an old version of data or no data at all under certain conditions. This compromise allows for maximum availability and scalability of the storage platform. Figure 4 illustrates the two types of consistency models – eventual and enforced. OpenStack Swift is an eventually consistent platform. IDC plans to use the CAP theorem in subsequent analysis of OBS platforms.

For a more detailed discussion of this taxonomy, refer to *IDC's Worldwide File- and Object-Based Storage Taxonomy* (IDC #239143, January 2013).

## FIGURE 4

---

### CAP Theorem



Source: IDC, 2013

## LEARN MORE

---

### Related Research

- *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)

## 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 1000 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For more than 48 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 2013 IDC. Reproduction is forbidden unless authorized. All rights reserved.

