

# Cool Vendors in Storage Technologies, 2017

**Published:** 12 May 2017 **ID:** G00327285

**Analyst(s):** Dave Russell, Arun Chandrasekaran, Julia Palmer, Raj Bala, Alan Dayley, Pushan Rinnen, Garth Landers

As I&O teams struggle to meet the demand for increased storage capacity, I&O leaders focused on storage will need to explore innovative technologies to support evolving business requirements. We highlight five disruptive private storage vendors for their abilities to address common pain points.

## Key Findings

- I&O leaders have identified "managing data growth to support business demands" as the top storage concern for 2017. Other major concerns include the need for high availability and performance of primary storage.
- I&O team capabilities have not kept pace with demands for increased storage capacity and rapid deployment, with only incremental improvements over the past three years.
- In addition to insufficient scale, a lack of automation continues to inhibit development.

## Recommendations

I&O leaders focused on infrastructure agility should:

- Reduce administrative demands by focusing not only on a new system's scalability, but also on its capacity for automation.
- Ensure that any deployment of automation is policy-based under the supervision of an I&O administrator to mitigate the requirement for deep storage expertise.
- Evaluate the following five disruptive private storage vendors based on their suitability to address specific workloads and use cases.

## Table of Contents

---

Analysis.....	2
What You Need to Know.....	2
Cohesity.....	4

- Elastifile..... 5
- Igneous Systems..... 6
- Komprise..... 8
- Portworx..... 9
- Where Are They Now?..... 10
  - Actifio..... 10
  - Druva..... 11
- Gartner Recommended Reading..... 12

List of Figures

- Figure 1. Top Storage Concerns for 2017..... 3

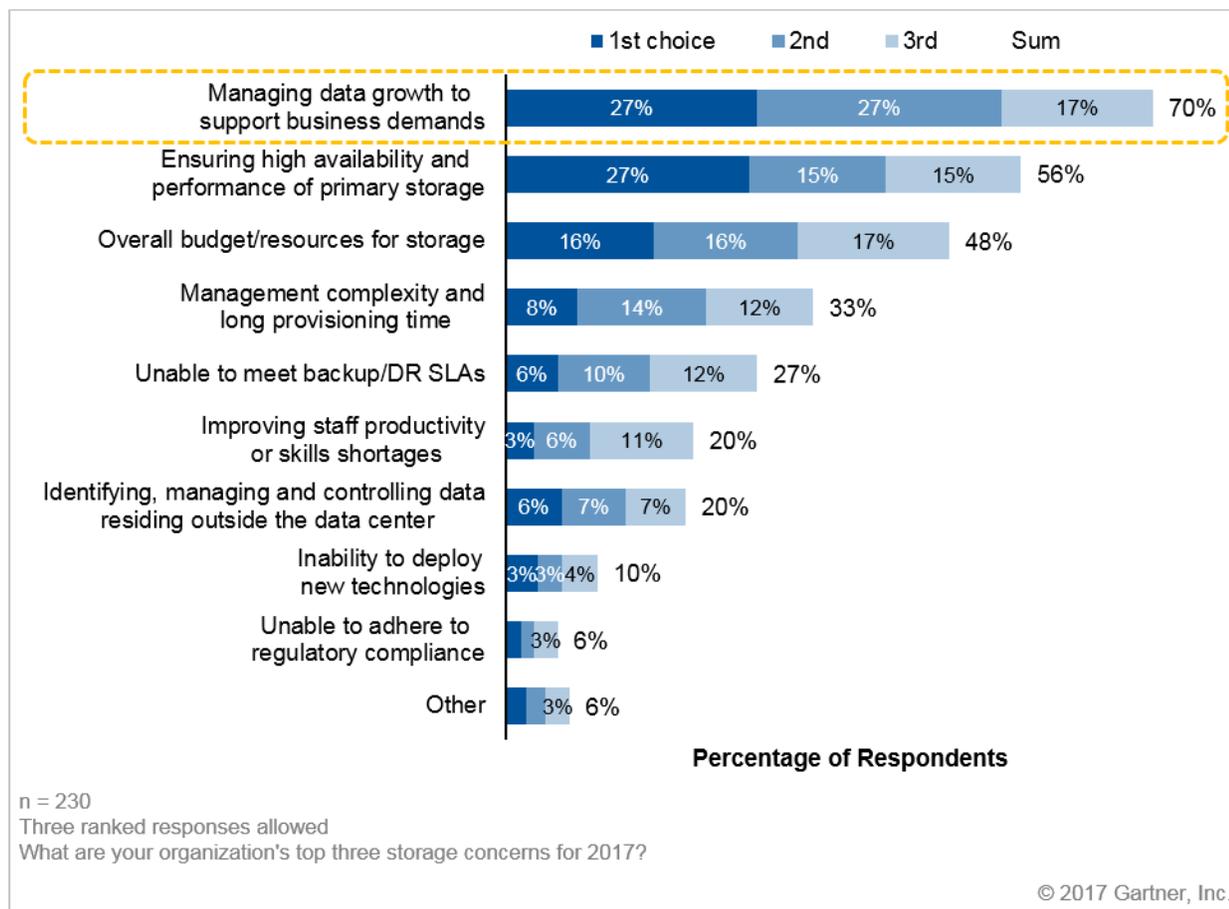
Analysis

*This research does not constitute an exhaustive list of vendors in any given technology area, but rather is designed to highlight interesting, new and innovative vendors, products and services. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.*

What You Need to Know

Infrastructure and operations (I&O) leaders often share common goals and challenges, including a need to modernize their storage infrastructures, improve agility and quality of service, and address the requirement to contain costs, while simultaneously delivering new applications and maintaining legacy systems (see Figure 1).

Figure 1. Top Storage Concerns for 2017



Source: Gartner (May 2017)

Managing data growth and supporting the business demands of providing storage have been the top concern of I&O leaders responsible for storage for the past three years. Sadly, current deployments and practices have done little to address this top pain point. The market is looking for solutions that can make significant, not incremental, gains in addressing what is becoming a more pronounced and more expensive issue.

It has become painfully evident that storage capacity demands, and expectations for far more rapid provisioning of that storage, have far outpaced the ability of I&O team's capabilities. Far-more-automated systems are required to restore a sense of balance, that is, storage solutions that offer much greater scale, but also much more automation. This required automation must be policy-based, and under the configuration, direction and ongoing supervision of an I&O administrator — ideally without requiring deep storage expertise.

We have chosen the five disruptive private storage vendors highlighted in this research because they exemplify innovative approaches to addressing the top storage pain point of managing data growth to support business demands. These vendors accomplish this through automated, scalable

storage solutions across a variety of workloads and use cases. These vendors can assist organizations in meeting their infrastructure agility, data center modernization and cost containment initiatives.

## Cohesity

Santa Clara, California ([www.cohesity.com](http://www.cohesity.com))

*Analysis by Dave Russell*

**Why Cool:** Cohesity is unique in that it is setting out to redefine the role of secondary storage. For decades, secondary storage has equated to archiving and backup data, as opposed to primary storage for active data. But Cohesity takes a broader view of secondary storage to include additional use cases, such as analytics and cloud-resident data, and also storage with less demanding SLAs, such as file shares, DevOps copies of applications and disaster recovery copies of data, in an attempt to automate the storage management of a very large number of workloads.

Cohesity envisions a world of high-performance, perhaps all-flash, primary storage, and then everything else. It is the "everything else," beyond the most demanding SLA workloads that Cohesity is setting out to manage. Cohesity would like to disrupt the current state of storage, with the many islands of fragmented secondary storage, which are typically encapsulated in several formats (i.e., that of the backup and the archives applications). Cohesity's hyperconverged secondary storage is architected to manage petabytes of data under its web-scale infrastructure that has a proprietary distributed file system (created by ex-Google engineers) and capacity-optimized global deduplication. This typically results in capital expenditure (capex) savings advantage. Cohesity also seeks to improve operating expenditure (opex) by collapsing the many tools, interfaces and policy settings that come with using multiple siloed solutions today (each of which requires its own management interface and stores data in a proprietary format), and replacing that with a unified policy engine and user experience.

Cohesity offers its DataPlatform solution via integrated, scale-out appliances for on-premises deployments, or via virtual machine (VM) instances in the cloud. DataPlatform can support connectivity simultaneously over Network File System (NFS), Server Message Block (SMB), and object storage (Amazon Simple Storage Service [S3] API). It also:

- Delivers in-line global deduplication, compression and encryption
- Provides unlimited snapshots and clones
- Offers write once, read many (WORM) retention on Cohesity volumes

Recently, it also delivered erasure coding to improve resilience to failures, while boosting usable storage over redundant array of independent disks (RAID) schemes by over 40%.

The optional DataProtect backup software protects workloads via subminute snapshots. Protection was initially focused on VMware, but now also includes physical server volumes (Linux, Windows), network-attached storage (NAS) filers, Pure Storage flash arrays, as well as physical Microsoft SQL Server and Oracle Databases (via Recovery Manager [RMAN] integration). Built-in analytics are offered for capacity utilization, performance, indexing and search, and custom queries via

MapReduce. DataPlatform can natively archive, tier and replicate data into cloud services (Amazon Web Services [AWS], Azure and Google Cloud Platform).

**Challenges:** The greatest issues for Cohesity are messaging, routes to market and rounding out its feature set. The messaging issue is largely around expanding upon what secondary storage is, as well as potentially having to prove itself to a larger audience with expanded roles in larger enterprises. The routes-to-market issue is squarely focused on growing its channel presence and the number of people that are involved in promoting the solution, and in which geographies. Cohesity needs to continue to round out its capabilities to add an iSCSI attachment and to protect broader workloads (e.g., Hyper-V and physical apps beyond SQL Server).

**Who Should Care:** Cohesity will appeal to a spectrum of I&O leaders. On the one hand, those looking to address the overall cost of backup storage compared to market leading deduplication appliances (the vendor claims typical savings of 50% or more) should investigate Cohesity for its industry-redefining view of what secondary storage is. This is also true for those looking to centralize management complexity of all but the most high-performance storage requirements under a single, global, scale-out, deploy and pay-as-you go solution. In the middle, those that seek lower cost, easier-to-use backup will want to examine Cohesity. Those seeking to better leverage their data, such as for disaster recovery or DevOps (via instant clones) for more rapid application development, will also be interested in Cohesity. At the other end of the spectrum, those looking for a more affordable, fully global deduplicated backup and/or archive solution for their current backup and archiving applications will find Cohesity appealing.

## Elastifile

---

Santa Clara, California ([www.elastifile.com](http://www.elastifile.com))

*Analysis by Julia Palmer and Arun Chandrasekaran*

**Why Cool:** Elastifile, a software-only startup, built a flash-optimized scale-out file system that can form a single namespace across on-premises and public cloud deployments. Elastifile has been awarded a patent for tiered storage in flash memory, where data is written to the appropriate tier based on the estimated write activity level. The technology supports a different mix of solid-state devices that might be available in the future. The solution features other patented techniques for dynamic data path optimization that can choose where specifically to write data. These have the ability to cancel and reselect write location in the event of errors, loss of the nodes or performance issues in large consolidated clustered environments with expected interruptions of node availability.

The Elastifile enterprise solution was designed to address the transactional NAS use cases where enterprises require the resiliency of distributed file systems without sacrificing performance at large scale for either high-throughput or low-latency workloads. The software can be deployed as a container or VM on the compute node as hyperconverged, or as an external file system based on a scale-out, shared-nothing architecture with distributed metadata management. Elastifile has been designed to scale to billions of files and thousands of nodes due to the fully distributed metadata and purpose-built key value consensus algorithm that enables high input/output operations per second (IOPS) and lower latencies across the cluster. In 4Q16, Elastifile released the first version of

its distributed file system software-only product for all-flash storage media called Elastic Cloud File System (ECFS). It supports both NFS and SMB file protocols and has global data deduplication and compression.

The Elastifile product was written with flash storage optimization features to deliver consistent high-performance file data services based on heterogeneous commodity server hardware of the customer's choice. As a result of venture investments from flash manufacturers, Elastifile has been able to work very closely with the flash vendors to minimize write amplification, and to gain visibility into usage and control of the underlying flash resources.

In April 2017, Elastifile added public cloud deployment support and CloudConnect tiering and migration features to enable enterprises to integrate their data management with public cloud offerings.

Elastifile pricing is based on a pay-as-you-grow, 100% subscription model, based on capacity consumption, and including enterprise support and services.

**Challenges:** With the first version of its product released in late 2016, Elastifile now joins the turbulent and crowded market of both established storage array vendors and emerging software-defined storage companies that are targeting existing data center shared file workloads. Although the Elastifile enterprise platform incorporates web-scale and cloud-like innovations, customers may be wary of investing in another siloed solution that will target only file-based workloads. In addition, end users would be required to provide solid-state drive (SSD)-based hardware. This may be hard to justify from an ROI standpoint given that only a minority of the existing enterprise file workloads require an all-flash storage tier. Elastifile's solution may be deemed both proprietary and expensive for common hybrid cloud storage use cases. While cloud builders and large enterprises are always looking for the solutions that can show web-scale efficiency with scalable performance, Elastifile would have to answer tough questions about roadmap, financial viability and total cost of ownership (TCO) against competitive on-premises and public cloud alternatives.

**Who Should Care:** Elastifile will appeal to I&O leaders and enterprise architects who are looking for elastic and scalable file services with a focus on predictable latency for transactional workloads. These workloads could feature databases, VMs and stateful containers, as well as high-performance computing requiring high throughput and scale. Cloud architects should explore Elastifile Cross-Cloud Data Fabric with its CloudConnect feature to deliver unified data services for performance-oriented use cases, such as analytics and the Internet of Things (IoT). Given its web-scale-inspired and cloud-ready nature, and its pay-as-you-grow consumption model, Elastifile could show better TCO to large enterprises or service providers that focus on hybrid cloud storage deployments of unstructured data with embedded unified management and data mobility.

## Igneous Systems

---

Seattle, Washington ([www.igneous.io](http://www.igneous.io))

*Analysis by Raj Bala*

**Why Cool:** Igneous Systems develops hardware infrastructure that is deployed on-premises, but managed by the company remotely for enterprises who pay for the product in a consumption-based model. An enterprise simply provides power and cooling while Igneous remotely manages everything else, including system upgrades and component failures. Upgrades occur, bugs get fixed, and new features simply appear without *any* effort on the part of the enterprise. This results in enterprises experiencing some of the most significant benefits of public cloud IaaS, but within their own data centers.

Igneous' product is a very early entrant in the microconverged integrated system (MCIS) space. This is an emerging architecture consisting of an Amazon S3-compatible object store and an AWS Lambda-like serverless computing environment that represents the convergence of compute and object storage on a single platform. The "micro" part of the MCIS acronym is represented by the Lambda-like serverless compute environment that executes a single function rather than a complete application.

As an S3-compatible object storage platform, enterprises can use Igneous for a wide variety of applications, including as a target for backup and archiving, as a data lake for big data applications such as Hadoop and Spark, and as a storage repository for cloud-native applications built by the enterprise.

Igneous' serverless forays include an environment by which functions, written by software developers, can automatically execute based on events that occur. This is very similar to AWS' Lambda public cloud service in which functions can be executed when a host of events occur on AWS. For example, storing an image using S3 can trigger a custom function that resizes the image for multiple screen resolutions.

Much like public cloud storage services, such as Amazon S3, Igneous' storage infrastructure is designed with the pragmatic sense that failures in a complex system will occur. Distributed systems such as Igneous are simply architected to handle failures better. In Igneous' case, every disk-drive-maintaining state has its own ARM processor card that performs I/O optimization and component failure management. The ARM processors attached to the drives can help rebuild data from erasure-coded fragments when drives fail, and perhaps, in the future, run more meaningful workloads that move compute operations closer to where data is stored.

Many of the Igneous management and engineering teams were instrumental in the success of EMC's Isilon scale-out storage product. As such, the people behind Igneous have experience in building infrastructure used by thousands of enterprises for workloads focused on unstructured data.

**Challenges:** While Igneous supports FTP and NFS (in read-only mode at the moment), its primary access mechanism is through the Amazon S3 API. There is a large ecosystem of software developers that have built tools and applications around the S3 API. However, this limits Igneous' appeal in an on-premises storage market where file protocols such as SMB and NFS continue to be dominant choices for unstructured data workloads in enterprise data centers. The MCIS category is a nascent architecture, consisting mostly of emerging vendors who are pairing a serverless compute environment with object storage, which is itself an emerging market with limited awareness outside

of public cloud IaaS. Moreover, serverless computing in enterprise data centers is so new that common use cases have yet to emerge, making monetization of the solution challenging today. As a result, Igneous is developing its own backup and archiving-focused applications based on its product's serverless compute capabilities to deliver functionality for common use cases.

**Who Should Care:** I&O leaders seeking public cloud benefits without perceived public cloud risks should evaluate Igneous. It will particularly appeal to I&O leaders who seek to lower storage TCO for unstructured data, and pay for it on a consumption basis while retaining stored data in-house. The lower storage TCO can be attributed to the hands-free nature of Igneous' system in terms of management, and the attractive pricing, which is currently \$0.015 per GB per month with a footprint starting at 212TB.

## Komprise

---

Campbell, California ([www.komprise.com](http://www.komprise.com))

*Analysis by Alan Dayley, Arun Chandrasekaran, Garth Landers*

**Why Cool:** What sets Komprise apart in a crowded file analysis market is the ability to effectively scale into multipetabyte environments and to efficiently remediate data location based on preference and cost. Komprise does not require agents or dedicated hardware. Instead, it scans hundreds of file types across many storage repositories, including NFS and SMB, during available nonpeak times via Komprise Observer deployed as a virtual machine. Observer VMs can be quickly added in order to manage petabyte-scale environments. In addition, scale is achieved by "storing" metadata gleaned during the scans to computational cubes instead of into a centralized SQL repository often used by competitors. Its metadata is fully distributed, and real-time mapping is maintained in a redundant manner on customers' persistent storage, which makes the Komprise engine largely stateless.

The offering differs from a traditional storage resource management (SRM) system due to the ability to manage any metadata, including personas gleaned from Active Directory, and to provide a cloud-based "director" to display the storage layout. It also differs by its ability to immediately provide remediation. Komprise provides an interactive analysis that shows "what if" scenarios before you actually migrate data to a different tier. The analysis is displayed in a user-friendly graphical format, and includes possible cost savings and storage reduction savings according to easily adjusted parameters. This allows storage administrators to fine-tune their storage tiers, maximizing savings, while at the same time taking into consideration required access service levels. Komprise also offers built-in high availability.

In addition to managing on-premises storage, Komprise has production customers that use it as a tiering engine, with cloud storage services including Amazon Web Services, Google Cloud Platform, Microsoft Azure and Oracle Storage Cloud Service. Komprise does not utilize stubs, but rather what it calls "dynamic links" to access the data, which shows up in the original directory with no detectable change to the end-user experience. If the data is fully migrated to a separate location, Komprise offers an interface for users to access the data in the new location. Komprise is available via capacity-based pricing — the amount of source data that is managed.

**Challenges:** Komprise serves a multitude of use cases, such as data analytics and governance, storage migration and cloud tiering. However, it faces competition for each of these use cases from both large established infrastructure vendors and innovative solutions from new market entrants. For storage migration projects, storage vendors often subsidize, or offer for free, their tooling; but these are proprietary and ephemeral, and often the data cannot be read directly from the new location. Cloud storage gateway vendors, as well as cloud service providers, offer storage devices that have in-depth functionality with broad support and close integration for on- and off-premises object storage solutions, but they result in new storage silos. As an early stage company, the depth of Komprise's reporting capabilities may not yet meet the exacting needs of enterprise customers that may choose best-of-breed tools from larger independent software vendors (ISVs).

**Who Should Care:** I&O leaders grappling with large volumes of unstructured file content will be intrigued with Komprise's ability to identify, report and migrate to alternative/less expensive tiers of storage, including cloud storage. Traditional file archiving has had a messy history that has included a variety of approaches such as hierarchical storage management (HSM) and the long-term retention of backup data. Komprise provides an intelligent means of moving data based on its value, while also providing transparent end-user access — something traditionally lacking in file analysis offerings. For enterprises that have been reticent about moving file content to cloud storage because of a lack of data management, Komprise can assist by identifying the business value (or lack thereof) of content before migrating it to IaaS environments. These reporting capabilities may also be useful to compliance and governance personnel; but thus far, Komprise has been more suitable for operational efficiency and cost reduction.

## Portworx

---

Los Altos, California (<https://portworx.com>)

*Analysis by Arun Chandrasekaran and Julia Palmer*

**Why Cool:** Portworx provides a software-defined storage product that turns commodity servers into a scale-out storage cluster, providing data persistence and data protection for containerized applications in an elastic and granular manner. It can also manage storage provided by cloud volumes or traditional storage arrays (SAN and NAS). The product line, PX-Series, is a distributed block storage layer that deploys itself as a container on every node in the cluster. It simplifies the setup and scaling of the cluster through its fingerprinting and resource discovery layer. This scans the nodes and coalesces available system resources with constant communication that ensures real-time dissemination of data on resources used, as well as the overall state of the cluster. The PX-Series has close built-in integration with popular container orchestration frameworks, such as Docker Swarm, Kubernetes, Rancher and Mesosphere DC/OS. Storage volumes can be provisioned with granular snapshot (frequency and retention), class of service, encryption and replication policies. The PX-Series is a unified storage product line that can provision either block storage or file-based storage, depending on workload requirements. The data is distributed across the cluster of nodes to ensure high availability with protection against disk and node failures. The metadata is distributed and stored along with the volume, thereby reducing any dependence or failure domains associated with centralized metadata servers.

The adoption of containers is rapidly growing in the enterprise and is starting to expand beyond the initial stateless service use cases. Portworx's container-centric data services, shared-nothing architecture and software-defined approach differentiate it in the marketplace when compared to other competitive solutions. Its software-led approach also provides flexibility to consume its product both in on-premises data centers as well as cloud IaaS — a fast-growing option for microservice applications. The PX-Series is available in two editions:

- PX-Developer, a free developer-centric offering with a three server, 1TB/volume limit
- PX-Enterprise, an enterprise-focused offering with a web management console, RESTful API for automation and 24/7 support

**Challenges:** While the adoption of containers has been viral among developers, enterprise IT is still cautious about their adoption due to the immature state of the market. It is also concerned with operational challenges around security/isolation, monitoring, networking, data persistence and management. Additionally, Portworx is an early entrant in a nascent market. This clearly poses revenue growth and customer acquisition challenges as the market needs to mature rapidly for Portworx to be a commercially viable entity. While Portworx's software-defined approach provides flexibility, many enterprise IT buyers still prefer "appliances" due to easier unified support and for faster deployment. Moreover, the competition will become intense as legacy storage vendors, cloud IaaS providers and Docker jostle for relevancy in this space.

**Who Should Care:** I&O leaders tasked with selecting and managing infrastructure for stateful container workloads should include Portworx in their shortlists for evaluation. DevOps leaders building platforms to run stateful applications in the cloud or on-premises looking for rapid deployment and automated operations can benefit from the elasticity, portability and container-native data services that Portworx can deliver.

## Where Are They Now?

---

### Actifio

Waltham, Massachusetts ([www.actifio.com](http://www.actifio.com))

*Analysis by Dave Russell*

Profiled in "Cool Vendors in Storage Technologies, 2012"

**Why Cool Then:** Actifio was founded in July 2009 by storage industry veteran Ash Ashutosh, who has several other successful startups to his name. The company initially focused on protecting what it referred to as "copy data" — a term that Actifio coined, which is now used by most other vendors (such as Dell EMC, Hewlett Packard Enterprise, IBM, Pure Storage, Veritas Technologies, and IT research firms) — as opposed to production or primary data. Copy data is a superset of backup, snapshot, replication, continuous data protection (CDP) and secondary instances of data for sharing or performing analytics on. It also includes longer-term archival data copies. The company's offering was initially referred to as a Protection and Availability Storage (PAS) appliance, with the lofty ambition of unifying what are typically siloed activities such as backup, disaster recovery, business

continuity, and test and development data. Additionally, Actifio positioned the product as being a facilitator of private or public cloud storage by virtualizing data and allowing it to be diverted to other locations (including the cloud) without the need for organizations to build out secondary data centers.

**Where They Are Now:** In February 2017, Actifio repositioned itself to align with how its solution is being used. Actifio leverages past capabilities of copy data virtualization, and has expanded its focus to now deliver what it calls Enterprise Data-as-a-Service (EDaaS). EDaaS still includes backup and disaster recovery, but now brings in agility and hybrid cloud mobility, along with resiliency. No longer a storage-specific solution, Actifio targets larger enterprises. It also targets very different buyers, such as the VP of infrastructure, application owners and data center executives charged with better exploiting the cloud. What was initially a product focused on tactically containing storage costs by reducing the number of extra copies of data, Actifio has shifted to being a data management platform that offers greater agility and speed of delivery.

Today, Actifio is used to manage over 250 petabytes of application data, which represents over 6 exabytes of logical data that resides on any vendors' storage. Actifio's 2,200 customers range from on-premises data centers to cloud service and managed service providers. Actifio has become profitable and claims that it has been issued or has filed for 51 patents.

Deployments today are focused less on backup alone and more on application delivery and DevOps, as well as multidirectional hybrid cloud environments. While the majority of the installs have been in North America and specific countries in Europe, Actifio has been recently gaining traction in more of Europe, as well as Japan.

**Who Should Care:** Actifio should appeal to a broad range of enterprise data center executives. Vice presidents of infrastructure seeking to radically redefine how data is managed and leveraged, both on-premises and in a variety of cloud environments, will appreciate Actifio for its ability to optimize the delivery of expanded data services. Also, those responsible for producing new applications with greater speed and in less time should investigate Actifio for its ability to reduce cycle time and provide greater application resiliency and portability.

## Druva

Sunnyvale, California ([www.druva.com](http://www.druva.com))

*Analysis by Pushan Rinnen*

Profiled in "Cool Vendors in Storage Technologies, 2012"

**Why Cool Then:** Druva launched an enterprise-level endpoint backup solution called inSync in 2010, uniquely leveraging AWS's native compute, database and storage resources. Initially, the product could scale to 10,000 endpoints per server, and offered source-side global deduplication across users for storage and network efficiency. Performance was driven by caching on both the client and server side to reduce storage I/Os over the network and by parallel backups. Compared to competitors, inSync was early to support mobile operating systems and security features such as

two-factor authentication, GPS tracking and remote wipe. It also had a modern-looking, intuitive user interface and dashboard with federated search.

**Where They Are Now:** Having raised \$118 million in five rounds of VC funding, Druva has evolved into a cloud backup and data management company with a much more comprehensive portfolio to protect endpoint devices, on-premises servers and cloud application data. Today it has two main products: InSync and Phoenix. InSync continues its path of unifying backup, recovery, security and compliance archiving capabilities for user data. It has added backup and archival for SaaS applications such as Box, Google G Suite, Microsoft Office 365 (with support for SharePoint Online planned in June 2017) and Salesforce. The product has expanded its use cases to include PC/laptop refresh/migration, data governance (legal hold and integration with e-discovery tools), and security monitoring and alerting for PII data compliance and ransomware. InSync now can also be deployed in Microsoft Azure, in addition to AWS and on-premises deployments. Most of its customers use cloud deployments now that such deployments can scale to over 100,000 devices for a single organization.

Druva's Phoenix product delivers converged public cloud backup, disaster recovery and archival for on-premises servers, including Windows, Linux, Microsoft SQL Server, and VMware virtual machines. The product is primarily deployed at remote and branch offices and small to midsize data centers. The optional CloudCache keeps 30 days of backup locally for performance, while global deduplication provides block-level, incremental forever backup and cloud network efficiency. Auto archival provides policy-based autotiering among AWS's different tiers of S3 and Glacier storage to lower cloud storage costs. Cloud backup copies of VMs can be mounted directly for disaster recovery or for test/development. Druva is planning to launch AWS storage snapshot management capabilities in June 2017.

**Who Should Care:** Druva's solutions appeal to organizations with "cloud first" strategies who are looking for solutions to deploy backup, recovery and archival, with no — or a drastically reduced — hardware footprint on-premises. This is especially the case for enterprises that want to replace tape backup in their branch or remote offices, where cloud connection is reliable, and network bandwidth is sufficient. Druva's backup as a service (BaaS) pricing model includes backup software and public cloud storage, without the need for organizations to procure and pay for their own cloud storage, simplifying ongoing management.

## Gartner Recommended Reading

*Some documents may not be available as part of your current Gartner subscription.*

"Infrastructure Agility Primer for 2017"

"2017 Strategic Roadmap for Storage"

"Predicts 2017: Business Continuity Management and IT Service Continuity Management"

"How to Determine Whether Software-Defined Storage Is Right for Your Organization"

"Large Established Storage Vendors, Brands and Products Are No Longer Risk-Free"

"The New Realities of Storage: The Age of 'Compressed Differentiation' Changes Buying Criteria"

### Evidence

This research is based on more than 3,000 client interactions conducted annually regarding how I&O leaders responsible for storage manage their storage infrastructure and processes. We also polled nearly 250 Gartner clients at the 4Q16 Gartner Data Center and Infrastructure and Operations Management Conferences in the U.S. and Europe on their backup and storage methodologies.

**GARTNER HEADQUARTERS****Corporate Headquarters**

56 Top Gallant Road  
Stamford, CT 06902-7700  
USA  
+1 203 964 0096

**Regional Headquarters**

AUSTRALIA  
BRAZIL  
JAPAN  
UNITED KINGDOM

For a complete list of worldwide locations,  
visit <http://www.gartner.com/technology/about.jsp>

---

© 2017 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. or its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. If you are authorized to access this publication, your use of it is subject to the [Usage Guidelines for Gartner Services](#) posted on gartner.com. The information contained in this publication has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This publication consists of the opinions of Gartner's research organization and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice. Although Gartner research may include a discussion of related legal issues, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner is a public company, and its shareholders may include firms and funds that have financial interests in entities covered in Gartner research. Gartner's Board of Directors may include senior managers of these firms or funds. Gartner research is produced independently by its research organization without input or influence from these firms, funds or their managers. For further information on the independence and integrity of Gartner research, see "[Guiding Principles on Independence and Objectivity](#)."