

Cool Vendors in Storage Technologies, 2016

Published: 22 April 2016

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

Emerging storage vendors offer data center managers and storage administrators fresh answers for their storage challenges. This research details several innovative storage vendors.

Key Findings

- Datrium enables easier storage management and greater performance by exploiting commodity, server-based solid-state drive resources, while offering more storage capacity than other server-based storage solutions.
- Hedvig provides unified access for bare-metal, hypervisor and container environments, along with enterprise data services, based on a software-only distributed storage platform that can be delivered in a hardware- and hypervisor-agnostic fashion, either on-premises or in the cloud.
- Minio is lowering the barrier to object storage adoption by making a product that's easy to deploy on a software developer's workstation and on servers in the data center.
- Peaxy provides seamless access to data and improves productivity, while aggregating multiple and disparate data sources.
- Rubrik offers a scale-out appliance integrated with backup software, a fast recovery storage tier and an onramp to a long-term-retention storage tier (local or in the cloud), as well as a global search tool for restore and recovery.

Recommendations

- Include Datrium as a potential storage solution for virtual-machine-centric use cases when decoupling capacity from performance in a cost-effective manner is a key priority.
- Consider Hedvig for a bimodal, unified storage platform for virtualized private cloud use cases where support for multiple protocols and delivery models (hyperscale/hyperconverged) is important.
- Use Minio when building or deploying applications compatible with Amazon Simple Storage Service. Experience Minio by first trying it on your Windows, Mac or Linux desktop computer by downloading a binary and running the executable.

- Examine Peaxy for use cases ranging from seamless access to data across disparate data sources to integrated search as part of a data access workflow.
- Select Rubrik for a simplified VMware backup infrastructure.

Table of Contents

Analysis.....	2
What You Need to Know.....	2
Datrium.....	3
Hedvig.....	4
Minio.....	5
Peaxy.....	6
Rubrik.....	7
Where Are They Now?.....	8
Delphix.....	8
Kaminario.....	10
Gartner Recommended Reading.....	10

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. New storage software and system providers can help stakeholders build easier-to-manage, more scalable and efficient storage infrastructures. Many organizations are evaluating technologies that will drive efficiency, such as higher-performing and more automated storage, cloud-based solutions, and data management and backup and recovery products that allow for proactive problem avoidance and increased resource utilization and optimization.

This research details five emerging private storage vendors that can assist organizations in meeting their infrastructure agility, data center modernization and cost containment initiatives.

Datrium

Sunnyvale, California (www.datrium.com)

Analysis by Arun Chandrasekaran and Dave Russell

Why Cool: Datrium is cool because its DVX storage system's architecture enables customers to leverage the performance and cost benefits of commodity server-side flash with the flexibility (ability to scale compute and storage independently) and the resiliency of back-end persistent storage. By leveraging the CPU cores on a group of clustered hosts for storage functions, the DVX architecture hopes to remove the controller bottleneck associated with traditional dual controller architectures, as well as the scaling issues that are encountered. The DVX product provides management at the virtual machine (VM) level, obviating the need to manage logical unit numbers (LUNs)/volumes, which makes it highly appealing to VM admins. Although DVX and hyperconverged integrated systems (HCIS) have a similar value proposition of managing and abstracting storage through VM management, DVX is differentiated by decoupling storage capacity from the compute capability to allow for more efficient utilization and expansion of shared storage resources, independent of server type and vendor.

Datrium provides an architecture that decouples storage capacity and performance by combining resources from the host with network-attached persistent back-end storage. DVX Hyperdriver software runs on any x86 VMware host and provides data services, such as deduplication, compression and cloning, while leveraging solid-state drive (SSD) or PCIe flash on the hosts to deliver a read cache. This read cache is local to each host, and the compute model can scale linearly by adding ESX hosts, with DVX supporting up to 32 hosts currently. Data is written synchronously to the battery-backed NVRAM in a Datrium DVX NetShelf appliance that is connected by a 10GE interface to the hosts, but a compressed copy is kept in the host cache for quick access, and this cache is deduplicated in-line. Datrium uses a proprietary protocol over Ethernet to communicate between the hosts and NetShelf. DVX NetShelf is a dual controller appliance with low-cost hard-disk drives (HDDs) that provide a cost-effective persistent data storage layer with RAID 6 protection computed on hosts. Writes are compressed on the host before being transmitted and written to the NVRAM tier. NetShelf provides global deduplication by leveraging CPU cycles from the ESX host for a postprocess deduplication.

Datrium was founded in 2012 by an executive team from Data Domain (EMC) and VMware, and its DVX software and NetShelf products became generally available in February 2016. DVX is priced based on a perpetual license model that is based on NetShelf's capacity, without any per-host fees.

Challenges: Datrium is competing in an extremely crowded market, with traditional hybrid storage arrays, as well as hyperconverged systems that can deliver many of these benefits. Although DVX's architecture is innovative, customers may be wary of the fact that this is not an integrated solution, as Datrium does not provide customer support for the host SSDs. NetShelf comes in a rigid 48TB single node configuration, with no scale-out capabilities at present. The product can offload clones of volumes, but lacks native snapshot and replication capabilities. However, Datrium recommends Zerto until those capabilities are built into the product, which has a target date of later this year. VMware ESXi is the only supported environment, with no support for either bare-metal, Hyper-V or

KVM hypervisors. DVX Hyperdriver software can only leverage PCIe flash or SSDs, and is unable to leverage DRAM on the servers for caching.

Who Should Care: I&O leaders focused on cost reduction can benefit from the ability to deploy commodity servers without sacrificing performance, and by deferring or avoiding expensive storage upgrades. Virtualization architects and storage architects looking to implement a high-performance and cost-effective storage solution for use cases such as virtual desktop infrastructure (VDI), virtual server infrastructure (VSI), and test and development environments should look at Datrium as a possible alternative approach.

Hedvig

Santa Clara, California (www.hedviginc.com)

Analysis by Julia Palmer and Dave Russell

Why Cool: Hedvig is cool because it delivers a complete unified platform with tunable data services to eliminate silos of protocol-specific vertical storage arrays. Hedvig designed its solution to provide storage flexibility and scalability for small to very large bimodal enterprises via its unified software-defined distributed storage solution.

Founded by experts in Web-scale development, the Hedvig Distributed Storage Platform is built as a true scale-out distributed system. The cornerstone of the architecture lies in a foundation independent of hardware, hypervisor, container or cloud platforms. Solution efficiency comes from the breadth of support for the underlying commodity hardware (x86 and ARM) and robust storage data services, where performance and reliability improve as the scale of deployment grows. Deployment can be in either a hyperscale fashion (to increase efficiency by decoupling storage and compute) or a hyperconverged manner (bringing storage closer to the compute layer for performance and data mobility), to deploy the solution on-premises or in the cloud. In addition to standard storage protocols and APIs — such as iSCSI, NFS, Amazon Simple Storage Service (S3), Swift and, soon, SMB — Hedvig provides Cinder and Mirantis Fuel drivers for OpenStack, ClusterHQ Flocker and drivers for Docker containers, and a VMware vCenter plug-in for VMware.

Customers can mix and match hardware configurations and storage policies to ensure that mixed workloads — such as database, Hadoop, VDI and VM, or development and testing (dev/test) — can all be run from the same platform. Hedvig supports major public cloud providers, enabling organizations to build hybrid storage clusters that stretch across multiple sites. The vendor provides administrators with granular fine-tuning of storage by enabling multiple storage and disaster recovery policies on a per-volume basis, which include disk size, disk type, disk residence (SSD, HDD or a combination), block size, single versus multireader, in-line deduplication, in-line compression, client-side caching, cluster-side caching, replication factor, replication destination, snapshots and clones.

Challenges: Hedvig Distributed Storage Platform delivered its first release in early 2015, and joined many software-defined storage vendors trying to gain market and mind share away from traditional storage solutions. While Hedvig promises Type A enterprises flexibility and freedom of choice for their hardware, hypervisor and storage delivery methods, overall implementations of software-

defined storage are still limited to early adopters. Although Hedvig's software-only approach is economically appealing for Mode 2 workloads, the selection, integration and rightsizing of the underlying hardware become end-user functions, which can contribute to more challenging deployments and integration, as well as inconsistent performance for some Mode 1 workloads. This issue could become more pronounced, because Hedvig does not limit itself to one workload, but rather promotes its solution as storage for a wide variety of use cases. Customers would have to work with Hedvig and hardware OEMs to optimize hardware, network and storage media choices to rightsize the cluster for mixed workload use. Hedvig will have to tell a compelling story on its total cost of ownership (TCO) against competitive on-premises and public cloud alternatives. Today, with a small number of large customers, the vendor will face questions regarding financial viability, solution scalability, reliability and performance at scale as it strives to win deals with cloud builders and large enterprise adopters.

Who Should Care: Hedvig will appeal to I&O leaders, enterprise architects and DevOps leaders, as well as to service providers looking for a scalable, multiprotocol, flexible storage platform for a bimodal data center and/or for consolidation and virtualization projects. Cloud architects should explore Hedvig for its ability to enable bidirectional hybrid cloud options, where data could dynamically reside on-premises, and in the public cloud, if needed, for mobility or cost containment. Infrastructure and DevOps leaders will find Hedvig's REST API attractive for its ability to provide self-service and data mobility for their business, while still supporting traditional enterprise storage needs for file, block and object. Given its scale-out and industry-standard hardware-based, software-defined nature, Hedvig could show better TCO to large enterprises that focus on petabyte deployments of unstructured data.

Minio

Palo Alto, California (<https://minio.io>)

Analysis by Raj Bala

Why Cool: Minio is open-source, Amazon S3-compatible object storage software developed with a microservice architecture in mind. Minio is written in the Go programming language, the main advantage of which over the alternatives is that it supports concurrency at the language level and is compiled to machine code. Minio is cool because it takes advantage of the concurrency aspects of Go and the scalability aspects of microservices, all freely available in binary and source code form.

Like Amazon S3, Minio is designed to be used for unstructured files, commonly referred to as objects. Minio is easily deployed on both a desktop operating system for a developer to use locally in the process of building an application and on server-class hardware when an organization deploys it for production usage.

Minio's issues and bugs are publicly managed on the repository's GitHub page. Anyone can submit a "pull request" that contains a bug fix, product enhancement or even an update to documentation. Minio recognizes that the complete user experience with a product is important even if the primary user is a software developer or storage administrator; thus, the aesthetics of the user interface is on par with well-designed consumer-grade services. Continuing with this philosophy of positive user

experience, Minio does not keep the management and administrative portions of its product as proprietary closed source, as is common with many commercial companies supporting open-source projects.

Challenges: Minio, like other object storage vendors, will find it challenging to convince developers that viable alternatives to Amazon S3 exist in the enterprise, particularly in isolation from the array of other services offered by Amazon Web Services (AWS). However, Minio being open source and having a very low barrier to adoption mitigates this challenge better than other object storage vendors. Additionally, developers are already conditioned to use Amazon S3 over the public Internet or a private object storage platform when building applications, because no suitable, local Amazon S3 runtime has been available in the past. As a result, developers may view Minio on their workstations as a solution to an insignificant problem. Minio also faces the challenge of remaining protocol-compatible with Amazon S3, which can be a moving target. Developers may build applications with Minio only to find that protocol behavior with Amazon S3 or other S3-compatible object storage platforms is different, resulting in poor developer experience. Minio is entering a market that is relatively small compared to other faster-growing parts of the storage industry. The market for enterprise-based object storage is growing at a 9.1% compound annual growth rate (CAGR), and overall awareness and adoption are low compared with newer categories of enterprise offerings, such as HCIS and solid-state arrays (SSAs).

Who Should Care: Storage administrators and IT leaders seeking better TCO, compared to block and file storage for unstructured data, should consider Minio for deployment on the server. DevOps leaders and developers who are building applications that utilize the Amazon S3 protocol can add Minio to their local development toolset, even if the organization uses some other S3 server, including Amazon itself, as the target. In other words, organizations need not use Minio on the server in production for application developers to get value out of it in their development environment.

Peaxy

San Jose, California (www.peaxy.net)

Analysis by Garth Landers and Alan Dayley

Why Cool: Peaxy was founded in 2012, and its Aureum platform is a software-based offering combining storage infrastructure and data management that aggregates a large volume of content from disparate sources under a unified namespace. Peaxy is cool because it provides high ROI by integrating data access in various workflows, including engineering specific ones where the product has gained early traction. Peaxy quickly and seamlessly enables access to indexed content regardless of location. This yields great productivity gains in environments where users previously struggled with searching, tagging and collecting from related, but siloed, content sets. The ROI that can be achieved through Peaxy is not trivial — one client determined that the technology had achieved \$6 million in productivity savings. The company does not present itself as offering a technology solution, but rather as being a business enabler. Peaxy has worked with buyers to provide a consultative approach, re-engineering workflows and business process improvement.

Peaxy solves several data management problems by unifying storage sources, including access rights, life cycle management and metadata management. The product uses Lucene indexing and supports custom parsers for named applications, such as AutoCAD. This enables search to be tightly integrated into workflows and enhances user productivity. Massive parallel processing ensures that nightly reindexing isn't required for successful search results. Role-based access policies are supported, regardless of where the content may reside and after it has been moved, and, more impressively, access policies can change based on the life cycle of the data. Peaxy consolidates datasets for access and analysis through an immutable namespace, wherever the data may reside, and this namespace is durable and can endure events such as migrations and storage refreshes.

Data management can also be event-related, such as during file share migrations or cloud enablement, and can also be used to enforce retention policies, purging data as it reaches expiration. Metadata management capabilities include details about data characteristics, such as data ownership, access rights and location, as well as insight into the value of the content through full text indexing where appropriate.

From a storage perspective, Peaxy supports petabyte-scale environments that can run on commodity hardware, supports standard file systems (including HDFS), and is Posix-compliant and hypervisor-agnostic. As part of its go-to-market, Peaxy is targeting Internet of Things (IoT) use cases, including ingestion and management of multiple data types (such as telemetry, simulation datasets, genomics, and sensor and streaming information). Early adopters include enterprises in aerospace, manufacturing, oil and gas, and intelligence agencies.

Challenges: Peaxy's sales cycles can be complex. While the end result and purchase of Peaxy offerings can be transformative, end-user buy-in and engagement can be a multistakeholder process spanning business roles such as engineering leader and process leader, often in multinational settings and those that involve business process change. These types of sales cycles can prevent growth by not allowing the vendor to scale the business.

Who Should Care: Industrial manufacturing environments with high volumes of unstructured content, like telemetry data; schematics; mission-critical, long-lived equipment; audio and video; and IoT data spread across fragmented heterogeneous data repositories, are good candidates for Peaxy. End-user use cases include data analysis, document search and look-up, remote access, and content collaboration.

Rubrik

Palo Alto, California (www.rubrik.com)

Analysis by Pushan Rinnen

Why Cool: In mid-2015, Rubrik launched an innovative, all-in-one, scale-out integrated backup appliance. It is cool because it can simplify enterprises' complex backup infrastructures by implementing agentless, policy-based backup management and backup storage tiering leveraging external low-cost storage, such as cloud, for long-term retention.

Although integrated (or, as Rubrik says, "converged") backup appliances are not new, early entrants typically don't have a scale-out server/storage hardware infrastructure, and many of them use their own traditional backup software that can still be complex to manage. On the backup software side, Rubrik applies in-line deduplication on the initial full backup and uses the incremental forever method afterward, with virtual synthetic full and continuous garbage collection in the background. Because it stores virtual copies of data in the native format, it allows instant read/write access for fast recovery. In addition, on the appliance, Rubrik maintains a global backup metadata index and search capability across the performance tier and the archive tier, which enables faster single-file discovery and recovery. Rubrik changes the traditional job-based management to SLA policy-based management, which simplifies the process of local backup and replication and file archiving, and can reduce the worker-hours needed to manage the backup infrastructure.

On the hardware side, the scale-out architecture with a distributed file system metadata service allows automatic expansion of compute and storage resources, with automatic failover and load balancing. Such a hardware architecture eliminates the difficult task of accurately forecasting the backup environment in three to five years, and avoids the common problem of undersizing the backup growth rate when purchasing new devices. The appliance functions as the performance tier for faster recovery, and it archives the old backups to a lower-cost tier of storage, which could be public cloud, object stores or NFS file storage.

Challenges: While Rubrik has created a scalable appliance technology foundation, its data management functions and support matrix are not as rich and wide as many backup software competitors. Rubrik currently supports VMware backup only, although it promises to support Microsoft SQL Server and Linux server in its next release in the first half of 2016. Meanwhile, more established backup software vendors are adding capabilities to make back-end storage more easily scalable. Similar backup architectures are also available from other startups. For back-end cloud archival storage, Rubrik only supports Amazon S3 and S3-compliant cloud storage, but does not yet support lower-cost cloud tiers such as Amazon S3 Infrequent Access and Glacier. The SLA-based (instead of job-based) management approach may appeal to IT generalists more than specialized backup administrators.

Who Should Care: Although Rubrik has deployments in large enterprises, it may be most relevant today for midsize organizations that run Windows and Linux applications on VMware and are looking for simple backup tools to reduce backup infrastructure complexity. Organizations with mixed physical and virtual environments may also find Rubrik attractive if they are willing to use a different backup tool for other hypervisors or physical environments — until Rubrik delivers support for these other platforms.

Where Are They Now?

Delphix

Menlo Park, California (www.delphix.com)

Analysis by Dave Russell and Garth Landers

Profiled in "Cool Vendors in Storage Technologies, 2013"

Why Cool Then: Delphix had been in business since 2008. In 2011, it began shipping its Agile Data Platform, a software solution that could significantly speed up application development, operations and the amount of time it takes to test copies of production databases. Reports, analytics and refreshed data for test and development environments went from days to minutes to become available. Delphix's database virtualization solution dramatically reduced the number of physical copies of a single database, thus significantly improving capacity utilization by decreasing the amount of storage and power consumed. Pointer redirection allowed for any point in time (APIT) recovery, and could use a single copy to rapidly serve up space-efficient full instances of database images for testing or recovery in minutes versus hours. At that time, Delphix claimed about 100 customers, along with 100 company employees, and offered support for predominantly Oracle Databases, with SQL Server just becoming available.

Where They Are Now: Since 2013, Delphix has doubled its customer base to over 200 typically very large enterprises and quadrupled its staff to more than 400 people. It has expanded the supported databases to include PostgreSQL, Sybase ASE, Microsoft SQL Server, MySQL and IBM DB2, as well as the full production environments for Oracle E-Business Suite, JD Edwards, Siebel, PeopleSoft and SAP ERP systems. Cloud deployments to AWS, IBM's SoftLayer and VMware Hybrid Cloud are supported. In April 2016, the company announced a CEO change, with the original CEO and founder transitioning to executive chairman and focusing on continued innovation in the Delphix platform.

In addition to application development and retirement, new use cases, such as Unix-to-Linux migrations, migrations to the public cloud, data archiving, and governance and compliance, including Dodd-Frank, are now activities for which Delphix is deployed. In support of this, Delphix acquired its partner for data masking, Axis Technology Software, in May 2015. In addition to these use cases, the combination of masking and data virtualization enables organizations to confidently move development and test environments to the cloud, and to offshore or outsourced partners. The company has delivered new products to enable end users to have a self-service ability to rewind an application to a past point in time, a managed service provider portal and an express limited version that offers a no-charge trial. Delphix now appears in two Gartner Magic Quadrants, one for structured data archiving and another for data masking.

The company doesn't typically sell to the traditional storage buyer. It claims the ability to release applications 10 times faster by getting the right data, to the right team, at the right time. The message resonates with CIOs, data center vice presidents and application vice presidents, as the ability to support development and testing, cloud migration, archiving and compliance, business intelligence, and historical analysis is a top IT initiative.

Who Should Care: CIOs looking to further their strategic IT initiatives, such as DevOps, migration to the cloud, and management of security and compliance, will be interested in Delphix for its ability to accelerate progress in these areas. Vice presidents of operations wanting to create new environments faster, vice presidents of applications looking to accelerate release schedules and vice presidents of infrastructure needing assistance with structured data archiving and application retirement/consolidation will want to consider Delphix. Service providers may want to examine the vendor for use in delivering improved infrastructure and services.

Kaminario

Needham, Massachusetts (www.kaminario.com)

Analysis by Stan Zaffos and Joe Unsworth

Profiled in "Cool Vendors in Storage Technologies, 2011"

Why Cool Then: Kaminario was founded in 2008 with the objective of building a no-compromise, high-performance scale-out storage system. The management team, headed by Dani Golan, obtained initial funding from Sequoia and Pitango, two well-known venture capital firms, and delivered its first SSA, the K2, a very fast (150,000 input/output operations per second [IOPS]) entry-level system that could scale to 1.5 million IOPS with 16-GB/sec of bandwidth in 2011. However, the focus on performance and bandwidth resulted in a system that was not cost-competitive with other vendors' SSA offerings. This was due to its use of expensive volatile DRAM technology, rather than much lower cost flash-based SSDs that still deliver a two-order-of-magnitude improvement in IOPS over HDDs, and its lack of compression and duplication — features that had already been promised by one or more of Kaminario's SSA competitors.

Where They Are Now: Kaminario obtained additional funding and rearchitected the K2 to efficiently implement compression and data deduplication, and to utilize the most cost-effective commodity SSDs to reduce costs without compromising availability, scalability or performance. K2 packaging enables K2 users to scale out and scale up to better align performance requirements with costs. Kaminario has pursued an aggressive marketing and sales strategy centered on reducing usable cost/GB, without compromising performance and availability. Although not yet profitable, Kaminario has gained a 3% share of the SSA market and receives consistently high ratings in Gartner's Critical Capabilities for SSA research (see "Critical Capabilities for Solid-State Arrays"). To grow revenue and better support its channel partners and customers, it has expanded its sales and marketing team, increased support staff, and relocated its headquarters to Needham, Massachusetts.

Who Should Care: IT leaders evaluating the deployment of SSAs in their data centers for high-performance, cost-effective application acceleration should consider Kaminario. Specific use cases include high-performance databases, virtual servers, VDI, business analytics and life sciences. Organizations that are dissatisfied with their current solution's scalability, performance, acquisition and ownership costs, and/or postsales service and support could benefit from inviting Kaminario to bid for their SSA business.

Gartner Recommended Reading

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

"I&O Leaders Can Benefit From Storage Industry Innovation"

"The Top Challenges Facing I&O Leaders in 2016"

"Infrastructure Agility Primer for 2016"

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

"Top Five Use Cases and Benefits of Software-Defined Storage"

"How to Leverage Public Cloud IaaS for Backup"

"Beware of Amazon S3 and OpenStack Swift API Compatibility"

"Storage and Data Management Are Not the Same Thing"

"The Top 10 Must-Ask Questions When Evaluating Solid-State Storage Arrays"

More on This Topic

This is part of an in-depth collection of research. See the collection:

- Cool Vendors 2016: Features, Fads and Disruptions Will Define the Digital Landscape

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>

© 2016 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](#)."