

- 1 Storing Data Without Knowing What's In It Is No Longer An Option
- 3 From the Gartner Files: Cool Vendors in Storage Technologies, 2015
- 10 Learn More About the DataGravity Difference

Transform Your Data from a Potential Liability to a Business Asset with Data-Aware Storage

Storing Data Without Knowing What's In It Is No Longer An Option

Managing Unstructured Information has become a Burden

Every day there is a new headline. Organizations, both large and small, are experiencing data breaches or leaks. It seems that being audited or exposed is nearly inevitable. Yet most companies maintain a “keep everything forever” policy due to potential business opportunities or regulatory mandates. Organizations are struggling to stay ahead of data growth rates, changing information management policies, and increasing business risks.

DataGravity, the industry's first data-aware storage, makes addressing these challenges possible. The DataGravity Discovery Series delivers best-of-breed storage with built-in data governance and security, search and discovery, and backup and recovery capabilities. Until now, bringing all of these capabilities together was cost prohibitive and complex. Managing the various solutions, such as a separate search tool or a separate backup and recovery application, demanded experts to implement and oversee the tools, thereby driving up costs and impacting storage performance.



DataGravity offers primary storage that can help you better manage, search and secure your data.

DataGravity provides companies with a holistic view of their unstructured data. Instead of the storage layer just being a repository for data, we see storage for what it needs to be: a strategic asset for your organization. While storage needs to deliver the capacity and performance that IT professionals have come to expect, we believe storage should do much more. Storage must provide insights into the data it contains and secure that information from those who intentionally and unintentionally put it at risk.



“Being listed as a Cool Vendor recognizes the disruption driven by the data-aware concept. We enable organizations to use the data they already own to increase business value and reduce risk.”

Paula Long
CEO & co-founder of
DataGravity

Understand Your Data at the Point of Storage

By combining state-of-the-art storage with zero impact protection, enhanced data governance and integrated search and discovery, the DataGravity Discovery Series system not only serves and protects data more simply, but also helps customers understand their data, bringing actionable insights and issues to the surface. Our intuitive UI helps IT professionals, storage, security and system administrators, and even line of business users, manage their data more effectively. DataGravity provides interactive reports and data visualizations that can expose insights, uncover potential risks and help more efficiently manage storage.

Architected to deliver and enrich insights on operations, people, and content at the point of data creation, data-aware storage provides a more comprehensive view of your information.

Gain Control of Your Data With DataGravity

With DataGravity, storage is no longer simply a repository for holding data. With data-aware storage, your organization can gain new levels of insight to more efficiently manage your data, stay ahead of the competition, and protect your valuable data while helping your bottom line.

The benefits of data-aware storage through DataGravity include:

- streamline data management capabilities
- gain better visibility into stored data
- identify and protect confidential and sensitive data
- safeguard data without impacting primary storage performance
- foster collaboration and empower end-users
- discover and protect files within virtual machines

Gartner, Inc. has included DataGravity in the list of “Cool Vendors” in the “Cool Vendors in Storage Technologies 2015” report. This report evaluates companies solving storage industry challenges with innovative solutions, capabilities and delivery methods that meet the needs of modern organizations.

Today, storage and IT administrators are increasingly expected to deliver data-driven IT strategy that impacts all areas of their business and holistically enhance team outputs, such as productivity, collaboration and security. Patented technology and a unique vision of storage as the epicenter of productivity, data security and competitiveness make DataGravity the natural choice for organizations looking to optimize the effective use and protection of their storage and data assets.

With DataGravity, data-aware storage delivers better data management through an intuitive user-interface that supports your ability to search and discover, protect your data and enhance governance and security - all at the point-of-storage. Now, organizations can achieve tangible business results based on data resources they were previously unable to tap into.

Included in this Newsletter is the full report “Cool Vendors in Storage Technologies 2015,” compliments of DataGravity. We invite you to learn more about how DataGravity and other emerging storage technologies are reinventing storage and empowering organizations just like yours to better manage, protect and secure your stored data.

The DataGravity Discovery Series presents data to help you better manage your storage. Our unique user interface surfaces insights, provides “at a glance” reporting and actionable information.



Source: DataGravity

From the Gartner Files

Cool Vendors in Storage Technologies, 2015

Emerging storage vendors offer data center managers and storage administrators new antidotes for their storage challenges. This research details five companies that provide innovative storage capabilities via new architecture and deployment methods, and looks back at two past Cool Vendors.

Key Findings

- DataGravity provides a primary storage platform that accentuates data awareness with built-in data governance, data protection, search and visualization capabilities.
- Infinidat provides a versatile, scale-out storage system with high availability and density while keeping the acquisition costs low.
- Infinio provides software for pooling server-side memory across multiple hosts to create a global, deduplicated acceleration layer that improves storage performance and initially requires no additional hardware.
- Maxta provides SDS that can be deployed on hardware of your choosing or is available through partners on preconfigured appliances to create a hyperconverged solution.
- Storiant is a multifaceted archiving platform that can support active, compliance and historical archives in a private cloud environment.

Recommendations

- Select DataGravity for workloads where data protection, governance and analytics are important business needs.
- Include Infinidat as a potential supplier in storage consolidation projects that can benefit from a highly reliable, energy-efficient, multiprotocol system at petabyte scale.
- Deploy Infinio when looking for a simple-to-use, nondisruptive means to increase the performance of your applications or virtual desktops, or to accelerate software build and test operations using existing server-side memory.

- Choose Maxta when compatibility with a variety of hypervisors and hardware is important and independent scaling of compute and storage is desired.
- Use Storiant when long-term, low-cost retention of petabyte scale data is needed.

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

IT managers often share common goals, including a need to modernize their storage infrastructures, improve quality of service and contain costs while simultaneously 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 input/output (I/O) products that allow for proactive problem avoidance and increased resource utilization and optimization.

This research details five emerging vendors that can assist organizations in meeting their storage modernization and cost containment initiatives.

DataGravity

Nashua, New Hampshire (www.datagravity.com)

Analysis by Alan Dayley and Arun Chandrasekaran

Why Cool: DataGravity provides a primary storage platform with built-in data awareness — a combination of data protection, data governance, file analytics, and search and discovery. The DataGravity Discovery Series arrays are dual-controller, hybrid arrays that support block (Internet Small Computer System Interface [iSCSI] only) and file protocols (Network File System [NFS] and Server

Message Block [SMB]), offer thin provisioning, compression and deduplication, and enable storage provisioning and cloning through a simple UI. Built-in mirroring enables the primary controller to focus on performance, while the replicated copy (served by the secondary controller) is used for data protection, data governance, search and discovery. The Discovery Series systems track and audit all interactions, identify sensitive data, and retain this information in a metadata repository. A clean UI provides a simple view to search and filter through data based on metadata of over 400 unstructured data types. In addition, data can be easily searched for sensitive information, including personally identifiable information (PII) stored in a full content index as part of an overall information governance and compliance strategy. Visualizations assist in drilling into specific datasets and enable further action, such as deletion, migration and quarantine of data. The secondary replicated copy also acts as a backup for data protection utilizing DiscoveryPoints that are similar to a snapshot while also allowing granular recovery.

With \$92 million raised in capital investment and over 100 employees, DataGravity has a strong financial base to leverage for accelerated growth. The first customer shipment of version 1.0 of the product was in October 2014.

Challenges: Historically, the acquisition and provisioning of storage and the analyses of the stored data have been done either by different groups within IT or in different sequences. DataGravity might face challenges selling to storage administrators who are looking for a high performance storage solution while also trying to engage data governance or storage architects that might be more interested in file analysis solutions at a different time during the business process. Most data governance personnel are more familiar with software solutions that provide metadata and content analysis used for governance and compliance objectives.

Who Should Care: Storage administrators, storage and IT architects, IT management, chief data officers, and information and governance leaders who seek a primary storage platform that accentuates data awareness with built-in data governance, data protection, search and visualization capabilities should consider DataGravity.

Infinidat

Needham, Massachusetts (www.infinidat.com)

Analysis by Arun Chandrasekaran and Stanley Zaffos

Why Cool: Infinidat's flagship product — InfiniBox — offers a highly scalable and dense storage solution with high performance and reliability and a lower acquisition cost than monolithic storage systems. InfiniBox architecture is a clean sheet design that targets the high end of the storage market. InfiniBox highlights include a claimed seven nines of availability, petabyte scalability, triple-controller mesh architecture with massive DRAM and flash cache for performance, block protocol support (with upcoming native network-attached storage [NAS] and object protocol support), low overhead snapshots, automated tiering, low disk failure rebuild times, integration with multiple hypervisors, and ease of management augmented by VMware and OpenStack support. InfiniBox supports Fibre Channel (FC), Ethernet and InfiniBand protocols.

While InfiniBox's technical specifications are impressive, it is Infinidat's promise of very low acquisition and ownership costs, tight user budgets, and the track record of the company's founders that is driving much of the market's current interest in the company. All software features are included at no additional charge in a simple licensing model. The vendor also claims an energy-efficient design where the power consumed is 8 kilowatt-hours (kWh) during peak load on a 42-unit (42U) rack.

Challenges: Infinidat was founded in 2010 and started shipping products in 2013. While the company is registered in the U.S., development is done primarily in Israel with senior marketing personnel split between the U.S. and Israel. Despite shipping products for the past few years, the company is unusually reticent on its product architecture, capabilities and future roadmaps. Moreover, little attention has been paid to marketing, which has resulted in low awareness about the vendor and the product.

For Infinidat to succeed in the high-end storage market, it will have to overcome the risk-averse behavior of the decision makers with tight software partnerships, provide enthusiastic customer references, and offer easy access to its

integration labs for proofs of concept (POCs) as well as flexible terms and conditions that do not devalue its systems.

Who Should Care: CIOs and vice presidents of infrastructure that are looking for a flexible, scalable and cost-effective storage platform should look at Infinidat as an instrument to drive better consolidation and operational simplicity in their storage environment. Storage administrators focused on high reliability and scalable performance with simplified management can benefit from InfiniBox's innovative architecture.

Infinio

Cambridge, Massachusetts (www.infinio.com)

Analysis by Dave Russell and Arun Chandrasekaran

Why Cool: Infinio's Accelerator product aggregates a user-configurable amount of RAM across a VMware ESX cluster and presents this pooled, shareable read cache across all servers to deliver accelerated performance for applications. Accelerator's read cache leverages global 4KB block-level deduplication that creates effective cache sizes that are typically three to six times the physical aggregated RAM and often much higher for virtual desktop workloads that have a great deal of redundant images. Other key attributes are that installation is very rapid, typically less than 30 minutes, and, as existing memory is utilized, no additional hardware purchases or installations are required. Quite unique to Infinio is the ease with which the product can turn host acceleration on or off with a single mouse click in the administrative console with no disruption or reboot required. The product can be quickly and nondisruptively removed, if desired. Organizations can seamlessly continue to use all the features on their shared-storage arrays (such as snapshots, replication, tiering and thin provisioning), and all VMware services (such as Dynamic Resource Scheduler [DRS] and vMotion) are supported and not altered in any way. In fact, Infinio is an early vendor that is working with VMware for integration with VMware vSphere APIs for I/O Filtering (VAIO) in vSphere 6.0.

Infinio's aim is to decouple storage performance from storage capacity by providing a server-side I/O optimization layer that brings the I/O closer to the application or virtual desktop workloads. The current product improves response time from storage by offloading reads to the server cache. Accelerator can be used to speed up existing "brownfield" infrastructures, possibly delaying

array upgrades or avoiding the purchase of unnecessary capacity, or as part of a new style "greenfield" storage architecture where server storage is leveraged for performance and shared storage is less relied upon for active data. The administrative console provides a drill-down of the I/O of the storage at a virtual machine (VM) level, making correlation of application and storage performance easy to view. Accelerator is installed as a set of VMs on an ESXi host with the virtual appliance using two vCPUs that, by default, leverage 8 GB of RAM per instance. The amount of RAM is adjustable, but Infinio claims that with its deduplication, large amounts of memory are not required to effectively offload 65% to 85% of read I/Os from storage arrays.

Infinio was established in 2011 and delivered version 1.0 of Accelerator in 4Q13. While 2014 was spent continuing product development and getting initial customers, goals in 2015 are to bring that development work to market and establish a reseller ecosystem.

Challenges: Infinio is a small, startup vendor trying to compete in a nascent, but fast-moving, I/O optimization market that has several emerging providers, larger NAND flash manufacturers (typically with acquired solutions), and potential alternatives from legacy storage providers as well as possibly hypervisor vendors themselves. The product currently offers only NFS support and is limited to VMware environments, restricting its total addressable market, although block protocol support is expected to be available in 2Q15. Adding solid-state and write-caching support would lead to a larger addressable market. Expanding awareness and routes to market will be vital as the product seeks broader adoption.

Who Should Care: Virtualization architects looking for a simple, nondisruptive means to increase the read performance of their virtualized applications or virtual desktops should consider Infinio. Storage architects looking to scale out input/output operations per second (IOPS), extend the life of their existing array or reduce costs on future storage purchases will also benefit from Accelerator. CIOs or vice presidents of infrastructure should look at Infinio as a component for new data center architectures that do not rely on large legacy storage arrays for delivering needed performance for new applications.

Maxta

Sunnyvale, California (www.maxta.com)

Analysis by Arun Chandrasekaran, George Weiss and Dave Russell

Why Cool: Maxta offers a software-defined storage (SDS) solution to build converged infrastructures that enable customers to create a VM-centric infrastructure environment with rich data services and easier scalability. When users are offered hyperconverged solutions, the vendor usually offers the user little choice in hardware configuration and virtualization. Realizing IT leaders are not uniform in their preferences and needs, Maxta allows customers to retain the choice of servers, storage devices, hypervisors (ESXi, Hyper-V and KVM) and containers through a software-only solution. For customers desiring a more turnkey solution, it partners with server platforms from Cisco, HP, Intel and Supermicro to offer an appliance. In addition to flexible choices and capex savings in software licenses, savings are accrued from storage capacity optimizations.

At the heart of the architecture is a distributed file system that is installed as a virtual appliance on all nodes and provides a global namespace for scaling out, which Maxta claims is optimized for random reads and writes of data. The Maxta software runs as a virtual appliance on ESXi environments but as a process in other hypervisors and containers. Storage block size is aligned with an application profile at VM creation to provide improved performance, and storage properties can be identified on a per-VM basis. The key data services delivered in the platform include thin provisioning, in-line data reduction (deduplication and compression), writeable snapshots, cloning, read/write caching, mirroring and remote replication. Availability and resiliency are offered via Maxta's own VM-level replication that operates across the nodes and that can sustain a triple device failure via a checksum history. Multiple copies of the data are written out and, in case of an error, an alternate copy can be fetched during the read process. Compute and storage can be scaled independently, allowing for deployment and growth flexibility. Consequently, Maxta addresses wide market opportunities through the channel, OEMs and service providers.

Challenges: Maxta has been shipping products since 2012 and competes with several startup firms and larger incumbent server/storage

vendors that have much larger financial backing and routes to market in the SDS and converged infrastructure space. While Maxta's software-only and hardware appliance approaches may be appealing to large service providers and Type A customers, the lack of a single throat to choke and its support team's bandwidth may be potential concerns. These solutions will demand growing Maxta's organizational maturity to manage fully integrated offerings at multiple layers of control and through generational life cycles of various OEM partners.

Who Should Care: Maxta will appeal strongly to service providers and large enterprises looking for choice, simplicity and life cycle support within a commodity-based software solution that can be used to build a hyperconvergence solution. I&O leaders who especially want to maximize their choices of server and storage components, hypervisors and/or containers retaining vendor-branded technologies will want to evaluate Maxta's version of software-only or appliances solution compared with other vendor appliance approaches.

Storiant

Boston, Massachusetts (www.storiant.com)

Analysis by Garth Landers and Sheila Childs

Why Cool: Storiant provides object storage, delivered at low cost and capable of managing petabyte scale datasets. The company was founded in 2012, is backed by two major venture capitalists and looks to grow through partners. The product is suited for use cases, such as long-term archiving (especially for large datasets, such as geospatial, seismic or machine-generated data) with an emphasis on support for long-term retention of data and analysis of that data through high-speed streaming APIs. Compliance-oriented capabilities include providing immutability and support for Securities and Exchange Commission (SEC) 17a-4 compliance requirements. The Storiant system uses low-cost commodity servers and storage, and guarantees 18 nines of reliability with two copies using a hybrid approach utilizing erasure coding and replication. The number of copies is configurable, as is the percentage of erasure code elements. Data is distributed across just a bunch of disks (JBOD) enclosures attached to storage servers, and any storage server can return any data object requested by a user.

The company embeds open-source components (OpenZFS file system, Ubuntu Linux, Apache Cassandra and NoSQL database for metadata) and supports open standards hardware as well. Dynamic array striping and the use of two independent 10 GB/sec network interfaces enable concurrent access to data and high throughput. The StoriAnt software controls power to drives separately and is aware of the physical position of individual drives within the physical enclosure, providing optimized striping, maximized data throughput and drive spin down, which StoriAnt claims uses 75% less power than traditional archival storage. In addition, StoriAnt frequently monitors and error checks drives to guard against bit rot. Access to data can be provided through a number of interfaces, including OpenStack Swift and Amazon S3 REST APIs, NFS/CIFS and support for Java-, Python- and .NET-based clients. The system is multitenant. The cost per GB is very low — sometimes at \$.01/GB/month, which is among the lowest in the industry.

Challenges: StoriAnt must develop a compelling rationale for adoption, beyond low cost and reliability for archive data, where it competes with tape-based systems, like IBM Spectrum Archive. The product's support for highly regulated environments and related use cases inoculates it against providers like Amazon in those areas, but, for the greater number of enterprises that are increasingly storing everything forever, Amazon will "check the box," based on factors like brand recognition and customer footprint. From a go-to-market perspective, StoriAnt must develop a larger portfolio of management service providers (MSPs) and partners to resell their solution.

Who Should Care: Enterprises in industries such as financial services, healthcare, energy and media/entertainment, where considerations like heavy regulatory burdens, very large content stores and long retention periods impact storage decisions, should consider StoriAnt. In addition, StoriAnt will appeal to IT directors and storage and application administrators seeking to provide business users with low-cost and durable access to analytical archives. The offering is well-suited for those looking for ultra-long retention of data, as StoriAnt can effectively double the life span of a drive by powering down. For those organizations that want the scalability of public cloud but prefer to house data on-premises in a private cloud infrastructure, StoriAnt is a viable option.

Where Are They Now? RainStor

Dayton, Ohio (www.teradata.com)

Analysis by Sheila Childs and Garth Landers

Profiled in "Cool Vendors in Storage Technologies, 2011"

Why Cool Then: RainStor offered a specialized database that provided extreme compression, support for retention management and fast retrieval of archived data. The company originated in the U.K., launched its first commercially available product in 2007 and expanded to the U.S. in 2009.

The product was targeted at very large (petabyte) scale archives, especially historical archives associated with online transaction processing (OLTP) systems, online analytical processing (OLAP) systems and machine-generated data applications, such as call detail records (CDRs). Compression ratios of 20-to-1 to more than 40-to-1 were possible through the product's patented compression, based on field and pattern deduplication, followed by algorithmic and byte-level compression. Rapid retrieval of data was supported across billions of records, via full SQL query and support for business intelligence (BI) tools, such as IBM Cognos. The product was well-received by organizations with extremely large volumes of data that needed data retention, WORM support and other compliance functionality.

The company offered its product almost exclusively through partners, including HP, EMC and Teradata. Its size, lack of brand recognition and broadly separated geographical locations (U.S. and U.K.) were challenges that made significant growth difficult.

Where They Are Now: In December 2014, the big data analytics and data warehousing appliances company Teradata acquired RainStor for an undisclosed sum. Teradata purchased RainStor to add compliant archival data store capabilities to its portfolio of products supporting enterprise Hadoop environments.

Just prior to the acquisition, RainStor reported more than 150 customers across the various platforms it supported. The company has spent the last few years extending and strengthening

its partnerships as well as enhancing support for Hadoop. New functionality includes enhanced retention management, tagging, legal hold at the record level, data masking, encryption and support for various WORM storage platforms, such as network-attached storage (NAS) and storage area network (SAN). RainStor now offers native support for the Hadoop Distributed File System (HDFS), enabling analytics on historical data directly from Hadoop clusters rather than having to move data to alternate platforms. Access to data is via MapReduce, HCatalog or SQL using metadata filtering to enhance performance. Prior to the acquisition by Teradata, RainStor had a strong relationship with EMC supporting integration of the product with Isilon.

The RainStor acquisition makes sense for Teradata, who, as a partner, had kick-started RainStor's development efforts for Hadoop. The company had many joint customers as a result of their partnership, and joint solutions include RainStor FastConnect for Teradata, a fast, bidirectional data transfer connector co-developed by RainStor and Teradata, and RainStor FastForward for Teradata, a product that migrated data from older Teradata backup and archive tapes to the RainStor repository. Teradata expects to continue to develop RainStor as a software-only product with a focus on supplying archive, retention and governance functionality for Hadoop, and will additionally position the product as a companion to the Teradata Aster Big Analytics Appliance and the Teradata Appliance for Hadoop. Development efforts will center on ongoing product enhancements as well as integration with the Teradata Unified Data Architecture.

Teradata has already defined a Teradata RainStor roadmap that is currently confidential. The RainStor team is operating business as usual and expects existing partnerships to continue. RainStor's installed base consists of large telco, retail, financial services and other institutions with large-scale data management requirements. The market for structured data archiving, application retirement and analytics archives, where RainStor plays, is largely untapped.

Who Should Care: Organizations with very large datasets (hundreds of terabytes to multiple petabytes) that need a historical or an analytical archive should consider RainStor. The products' ability to complement existing relational database management system (RDBMS), data warehouse

and Hadoop environments with functionality that enables compliance, compression and fast access to historical data provides significant cost benefit. If initial commentary coming from Teradata about the acquisition holds, potential customers may consider RainStor a safe bet for non-Teradata environments. However, organizations that own products that compete with Teradata solutions should watch roadmaps carefully. Certainly, organizations that are customers of Teradata, particularly the Teradata big data products, should take a close look at RainStor.

Pure Storage

Mountain View, California (www.purestorage.com)

Analysis by John Monroe and Joseph Unsworth

Profiled in "Cool Vendors in Storage Technologies, 2013"

Why Cool Then: Pure Storage offered a compelling solid-state array (SSA) predicated on flash-focused data services, particularly in-line deduplication and compression, while leveraging inexpensive, consumer-grade solid-state drives (SSDs). Pure's array targeted highly virtualized environments, such as virtual desktops and virtual servers, but it also was gaining traction in database environments as end users began to address multiple workloads. Clever marketing campaigns helped propel its simple message and end-user-friendly business models to the industry and throughout its channel partners as it pushed throughout the U.S. Pure raised over \$100 million in its series D funding in 2012, and it had over 200 customer wins amid a very crowded and noisy landscape of both emerging and established vendors bringing their own solutions to market.

Where They Are Now: Since 2013, Pure Storage has witnessed a very strong growth trajectory in terms of revenue, customers and personnel. While Pure has enhanced its solid-state array products by upgrading its hardware to achieve greater capacity and performance, it also has added most of the remaining data management features to provide a more complete SSA. Pure's strength remains in its data management software, and, while usage of cost-effective consumer SSDs has been effective, it has become an impediment to providing more dense arrays and, thus, remains a weakness in the architecture. More importantly, Pure has continued to appeal to a wide range of customers through creative marketing campaigns that address the major pain

points of customers buying traditional storage arrays. All-inclusive pricing, maintenance protection and simple service-level agreements have resonated well with customers and aided Pure's continued expansion via its channel partner program. The company's go-to-market execution and its ability to scale its support and service while maintaining favorable customer satisfaction have been important testaments to Pure's strategy and culture.

Pure has also addressed long-term viability concerns by raising over \$470 million in funding and engaging in a cross-licensing agreement with IBM to address litigation from EMC. As a result, Pure has been able to attract world-class talent and grow its head count by nearly eight times since 2013, helping fuel its global expansion. As with nearly all vendors in the nascent solid-state

array industry, Pure is still striving to achieve profitability, which will become particularly important as it pursues a path to an initial public offering (IPO).

Who Should Care: Organizations that require high performance and availability across multiple application workloads in highly virtualized environments should consider Pure Storage. Pure's technology is a strong contender in the solid-state array marketplace and it should be considered a viable company with easy-to-use products and competitive pricing models for any organization pursuing a dual-vendor strategy or considering an emerging vendor.

Source: Gartner Research G00250240, Dave Russell, Arun Chandrasekaran, Sheila Childs, Garth Landers, Alan Dayley, Stanley Zaffos, George Weiss, John Monroe, Joseph Unsworth, 22 April 2013

Learn More About the DataGravity Difference

Are You Ready to Be Data-Aware?

DataGravity is being used by forward-thinking organizations to solve the most challenging unstructured data issues they have.

Interested in learning more? Read the [use cases](#) available on our Website to see examples of data-aware storage in action and learn how DataGravity can help your organization:

- manage confidential and sensitive data
- search and discover your unstructured data
- support regulatory compliance
- protect and rapidly recover your data
- identify dormant data
- see who is accessing your data

If you'd like more detail about our unique technology and patented architecture that makes the DataGravity Discovery Series the premiere storage solution for small and mid-sized businesses and organizations [download our white paper](#) on the technology that makes data-aware storage possible.



About DataGravity

DataGravity helps organizations unlock the value of their data. The DataGravity Discovery Series is a data-aware storage platform that allows IT professionals and line-of-business users to store, protect, search and govern their data. Its patent-pending architecture, powered by the DataGravity Engine, analyzes data as it is ingested without impacting performance, so administrators and users can quickly and easily explore and use data more effectively to derive insights that increase productivity, efficiency and organizational success. Headquartered in Nashua, N.H., DataGravity is led by visionaries in the fields of storage, business intelligence, data management and analytics and is funded by leading investors, including Accel Partners, Andreessen Horowitz, CRV and General Catalyst Partners. For more information, visit www.datagravity.com.

Source: DataGravity
