

Cool Vendors in Storage Technologies, 2014

Published: 15 April 2014

Analyst(s): Dave Russell, Arun Chandrasekaran, Gene Ruth, Pushan Rinnen, Stanley Zaffos

Emerging storage vendors offer data center managers and storage administrators fresh answers for their storage challenges. This research details four companies that provide innovative storage capabilities via new delivery methods and performance management, and looks back at two past Cool Vendors.

Key Findings

- Conduktiv Technologies offers input/output (I/O) optimization and acceleration that improves Windows application performance and can overcome the "I/O blender" effect of server virtualization, thus improving storage response times and containing costs.
- Exablox brings a simple-to-use, feature-rich, high-availability, scale-out network-attached storage (NAS) solution to small or midsize businesses (SMBs) using its own intellectual property and a "bring your own disk" strategy that disruptively lowers acquisition and ownership costs.
- PernixData provides virtualization software that can aggregate server-side flash devices across multiple hosts to deliver accelerated read/write performance with fault tolerance.
- Zadara Storage offers storage systems priced on a usage basis, and challenges well-established, capital-intensive storage hardware purchasing patterns.

Recommendations

- Include Conduktiv Technologies' V-locity when implementing server virtualization that includes many Windows virtual machines (VMs) and/or whenever Windows application performance needs to be improved.
- Place Exablox in environments that require Common Internet File System (CIFS)/Server Message Block support and value deduplication, nondisruptive capacity upgrades, consistent performance, and the possibility of redeploying already-purchased disks.
- Deploy PernixData when looking for simple, nondisruptive means to increase the read/write performance of your virtualized applications using server-side flash.

- Include Zadara Storage where third-party-supplied (public and private) cloud block storage is desired to meet demanding performance and availability requirements.

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 and improve quality of service, while simultaneously containing costs. New storage components, solutions and management tools can help stakeholders build easier-to-manage, more efficient and more available storage infrastructures. Many organizations are evaluating technologies that will drive efficiency, such as higher-performing and more automated storage, cloud-based solutions, and/or I/O products that allow for proactive problem avoidance and increased resource utilization and optimization.

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

Conduktiv Technologies

Burbank, California (www.conduktiv.com)

Analysis by Dave Russell

Why Cool: Conduktiv Technologies offers several Windows-oriented software products; however, it is the V-locity I/O optimization offering that garnered Conduktiv a spot in this year's Cool Vendors research. V-locity is pitched with the tag line "50% performance improvement with no new hardware required." The vendor was founded in 1981 and eventually was called Diskkeeper, perhaps best known for its defragmentation product by the same name. In 2012, the company changed its name to Conduktiv Technologies, to better reflect its focus on nondefrag offerings, as well as the shift from commercial to corporate and OEM revenue streams.

V-locity leverages the vendor's deep Windows I/O knowledge to accelerate performance. While V-locity can also improve applications in a physical environment, the majority of the nearly 2,000 customers use the product to reduce the I/O blender effect in a server-virtualized environment, where many VMs running on a single physical server make the I/O traffic appear very random. This randomness is caused by reads and writes that are being intermixed across several VMs, which results in slower overall performance. To combat this, storage, virtualization and/or application administrators often overprovision storage to use more disk spindles and use wide striping to leverage more disks. This results in lower storage utilization and higher costs, and potentially

requires more space to account for more disks, not to mention the potential for more power and cooling requirements as a result of the extra disks.

V-locity optimizes reads and writes at the source, reducing the I/O required for any given workload. This is achieved by packing more data in each write request, and caching hot data in available server memory. The product overcomes the inefficient write I/Os of the Windows operating system by forcing writes to be written sequentially and preventing Windows split I/Os, whereby files are split across multiple I/Os. I/O read bottlenecks and delays are minimized with a self-learning, application-aware capability to precache data.

Challenges: While I/O issues are now better understood, especially in a server-virtualized environment, awareness of the best way to address them remains muddled. Conduvis will have to find a way to break through the fray of Peripheral Component Interconnect Express (PCIe) flash card and hybrid storage array vendors and get its message across a very loud landscape that is filled with larger, more established vendors. While V-locity is complimentary to other solutions, such as Fusion-io's PCIe flash and VMware's vSphere Flash Read Cache, Virsto file system and Virtual SAN — and can also operate and enhance any storage vendors' Windows-compatible disk arrays — finding buyers and then helping them understand that even these other investments can benefit from V-locity could be a complex task. Today, the solution is limited to Windows environments only.

Who Should Care: Storage and/or virtualization administrators looking to increase I/O without overprovisioning storage or necessarily having to deploy PCIe flash cards should consider Conduvis Technologies, as should virtualization administrators looking to reduce the I/O blender effect of stacking many VMs on the same physical server. Windows database (Oracle and SQL Server) administrators looking to maximize the performance of their applications should also consider the vendor, as should storage managers looking to garner more I/O operations per second (IOPs) and, therefore, more useful life from their existing storage arrays.

Exablox

Sunnyvale, California (www.exablox.com)

Analysis by Stanley Zaffos

Why Cool: Exablox has rethought the design, delivery and support of storage solutions targeting the needs of SMBs with limited IT skills and budgets in creating its OneBlox storage system and designing a support infrastructure. Primary OneBlox design objectives included a new simplified design, low unburdened manufacturing costs, simple configuration and ease of management, and efficient remote troubleshooting. OneSystem, a cloud-based management tool, eliminates the need for an on-site service console, and is providing opportunities for improving maintenance effectiveness by easing the consolidated collection of performance data and failure signatures. As this telemetry data grows over time, Exablox has the potential to improve maintenance effectiveness and provide users with proactive configuration guidance.

By developing a cloud-based management tool that eliminates the need for an on-site service console, Exablox has reduced product costs and configuration complexity. The vendor has also

simplified remote troubleshooting. These two capabilities can lower maintenance costs while providing opportunities for improving maintenance effectiveness insights over time.

Exablox is offering users confronting greater storage demands than budget resources a potentially attractive alternative to cloud storage, as well as the option of buying from nontraditional storage vendors (such as software and networking vendors seeking new revenue opportunities), or building their own storage systems using open-source designs. The alternative for users is buying OneBlox storage systems and support from Exablox, and buying their own storage (hard-disk drives [HDDs] or solid-state drives [SSDs]) from distributors or local computer superstores without the two- to six-times markups that most traditional system vendors charge their customers. Because disks can easily account for more than 50% of a disk storage system's hardware costs, and maintenance spend rates are tied to hardware and software spend rates with the vendors, the bottom-line savings of the Exablox approach can be in the order of magnitude (10x) range. OneBlox also offers users the ability to add or swap out capacity from already-installed OneBlox systems, and to add additional OneBlox systems to the OneBlox infrastructure nondisruptively, with zero configuration changes.

Challenges: OneBlox's lack of Network File System (NFS), VMware and Hyper-V support and market-validated performance benchmarks limit Exablox's target market primarily to capacity and file sharing applications, and backup/recovery. The lack of VMware and Hyper-V support may also make users somewhat more reluctant to do proofs of concept (POCs) because of the need to dedicate physical server resources to the POC. Like other emerging storage companies, Exablox will have to market-validate its product, support capabilities, and build and maintain sales momentum in a very competitive market. It will also have to demonstrate that its use of a distributed object storage back end does not limit the usable scalability of OneBlox storage systems.

Who Should Care: SMBs, healthcare, education, nonprofits, government agencies and managed storage providers fit within Exablox's target market for a variety of overlapping reasons and use cases. For SMBs and managed storage providers, they are ease of use, nondisruptive capacity upgrades and low cost per GB; for healthcare, education and nonprofits, it is low acquisition and ownership costs; and for government agencies, it is managing the RFP process to avoid having to install the "wrong" low-bid storage solution win.

PernixData

San Jose, California (www.pernixdata.com)

Analysis by Arun Chandrasekaran and Dave Russell

Why Cool: PernixData's flagship product, FVP, is hypervisor-based software that aggregates server flash across distributed server hosts to deliver accelerated read and write performance to existing shared storage.

FVP allows customers to decouple storage performance from capacity by leveraging a distributed, high-performance flash tier close to the applications. The fact that the product supports clustering of flash across multiple hosts means that VMs on a specific host can access flash devices placed on other hosts seamlessly. Since there is an aggregated flash layer that is shared by all the hosts,

reads from the disks are significantly reduced, even when VMs migrate to a new host. The FVP software supports all servers and flash devices (PCIe and SSDs) on VMware's hardware compatibility list, and works with all block storage today (with file storage on the road map). Early customer references have indicated that the deployments are nondisruptive, simple and easy to configure (parameters such as cache and replication settings are the only settings to adjust).

FVP installation is nondisruptive to customer environments, requiring no changes to the VMs or applications, and without the need for any agents. It is easily installed inside a VMware ESXi hypervisor and provides support for VMware services (such as vMotion, Distributed Resource Scheduler [DRS] and vCenter Site Recover Manager [SRM]). Unlike most other products in this space, FVP is capable of not only read acceleration (with write-through to disk), but also write-back acceleration. In addition, writes are synchronously replicated across clustered hosts to ensure high availability. There are architectural differences that enable clustering, write-back acceleration, fault tolerance and dynamic resource scheduling, which differentiate FVP from other solutions.

PernixData was established in February 2012 and came out of stealth mode in August 2013, when FVP 1.0 was launched.

Challenges: PernixData is a small startup vendor trying to compete with server-side flash software from established vendors, such as EMC, Fusion-io, SanDisk and VMware. While clustering and write-back acceleration differentiate FVP from VMware's vSphere Flash Read Cache, PernixData needs to continue to differentiate and diversify its hypervisor support. Currently, FVP's hypervisor support is limited to only VMware ESXi (versions 5.0.x, 5.1.x, 5.5.x), with no Hyper-V support. The product currently supports block-level protocols (FC, FCoE and Internet Small Computer System Interface [iSCSI]), with support expected for NFS soon.

Who Should Care: Virtualization architects and application owners looking for simple, nondisruptive means to increase the read and write performance of their virtualized applications should consider PernixData. Storage teams looking to scale out IOPS, extend the life of their existing array or reduce costs on future storage purchases will also benefit from FVP. By virtualizing the server-side flash tier with built-in fault tolerance, IT teams can deliver better performance and high availability, yet reduce the costs associated with more-complex upgrades.

Zadara Storage

Irvine, California (www.zadara.com)

Analysis by Gene Ruth

Why Cool: Zadara Storage presents storage as if it were an on-premises NAS or storage area network (SAN), when, in fact, the storage is cloud-based. Unlike traditional enterprise storage, the vendor follows the cloud model by using subscription pricing based on SLA/quality of service measures, rather than requiring an outright hardware/software acquisition. Founded in 2011, Zadara Storage offers enterprise block and file storage as a service. The vendor delivers the storage service as virtual private storage within existing public clouds, or as a turnkey solution delivered on-premises as a private or hybrid cloud solution. Current partners include Amazon Web Services

(AWS), Dimension Data, CloudSigma, KVH (Japan), SerenITaaS (Australia and Canada), CoreSite, Vivavo (Singapore), Equinix (globally) and Toshiba (Japan). Hardware and software partners include Cisco, NetApp, Huawei, IBM, Red Hat and VMware. The solution offers a means for organizations to move from a capital expenditure (capex) to an operating expenditure (opex) model, and avoid the administrative overhead of managing their own storage hardware environments and dealing with equipment upgrades and migration.

Zadara provides organizations with a high-performance alternative to native public cloud storage. For example, as compared to Amazon's Elastic Block Store (EBS) storage, Zadara provides high-performance, shareable block volumes; large volume size; support for NFS and CIFS; instant and cloneable snapshots; and at-rest encryption with keys held by the customer. In addition, tight control over the physical disk drives for wiping is available if needed. Snapshotting is interesting for those that require database snapshots for disaster recovery purposes.

Challenges: Offering a cloud-based storage purchasing model means that Zadara must overcome the entrenched mentality of outright hardware purchases that predominates traditional IT organizations. Potential customers must be forward-leaning and amenable to cloud storage environments, whether private or public, that also necessarily include cloud compute capabilities to complement the storage. Since Zadara is a new and relatively small company (approximately 100 customers, to date), IT organizations must be prepared to accept the uncertainty and foibles of a new vendor, including a new purchasing approach to storage. Although Zadara has, in some cases, partnered with large storage vendors (e.g., NetApp-Amazon and Toshiba [in Japan]) to build the underlying platform, the entire Zadara platform has not yet hardened via usage by a large number of customers over a long period of time.

Who Should Care: Zadara storage is attractive to public cloud users looking for higher performance and more resilience in a cloud compute environment. As a cloud-oriented offering, Zadara storage should be interesting to cloud services providers and managed service providers that wish to rebrand a premium block or file storage service or provide a colocation service. For customers wishing to keep storage on-premises (e.g., government agencies and media companies), Zadara provides the path to an on-premises, opex financial model, without the necessity of putting sensitive data and applications in a public cloud environment. NetApp customers can extend their disaster recovery into the Amazon infrastructure (via Zadara's elastic NetApp Private Storage for AWS) by using Data Ontap's SnapMirror feature directed at the Amazon-offered Zadara storage service, which can also migrate data into Amazon Simple Storage Service (S3) as desired.

Where Are They Now?

Isilon Systems

Hopkinton, Massachusetts (www.emc.com)

Analysis by Arun Chandrasekaran and Pushan Rinnen

Profiled in "Cool Vendors in Storage Systems, 2007"

Why Cool Then: Isilon Systems offered a clustering solution for scalable, but affordable, unstructured content storage with NAS access protocols. Among various symmetric cluster file systems for scalable performance on the market, the vendor stood out with its easy-to-deploy clustered storage appliance approach and well-rounded feature sets that targeted digital media and enterprise content storage. Highlights included a tightly integrated file system, volume manager and redundant array of independent disks (RAID) in one software layer; a clusterwide snapshot at a granular level; asynchronous replication; high availability with multiple failed nodes; fast disk rebuild time; and a policy-based migration tool.

Isilon was achieving success in a number of vertical industries that required large content storage, such as media/entertainment, Web services, cable/telecom, healthcare/life sciences, and oil and gas.

Where Are They Now: After it had its first profitable quarter in early 2010, Isilon was acquired by EMC for \$2.25 billion in the fourth quarter of 2010. Isilon's revenue was close to \$250 million in 2010. Since the acquisition, EMC has made significant investments in Isilon, growing its number of employees from 450 in 2010 to more than 1,300 today. Revenue has quadrupled in the past three years to more than \$1 billion in 2013, which is an indication of how successful the acquisition has been for EMC. Today, Isilon has more than 1,100 customers and is fully integrated into the EMC Velocity channel program, which has significantly helped its international expansion.

While media/entertainment, healthcare/life sciences, and oil and gas still account for a major portion of its revenue, Isilon has made significant inroads in the financial services and education/research sectors recently. New use cases (such as video surveillance and analytics) have been growing rapidly as well.

In its OS refresh in late 2012 (OneFS v.7.0), Isilon added several capabilities, such as file-level writable clones, authentication zones, VMware vStorage APIs for Array Integration (VAAI) and VMware vStorage APIs for Storage Awareness (VASA) support, and enhanced write once, read many (WORM) functionality. Since then, it has also added block-level deduplication and integration with Hadoop, Syncplicity (for file sync and share) and EMC ViPR. Due to the cumulative efforts mentioned above, Isilon earned a product viability rating of Outstanding, the highest rating awarded, in Gartner's Critical Capabilities research for scale-out file system storage (see "Critical Capabilities for Scale-Out File System Storage").

Despite this tremendous growth, Isilon still faces several challenges. Geographically distributed deployments of Isilon are rare due to the replication overhead and lack of dispersed erasure coding. Backing up a large amount of Isilon storage continues to be a challenge for customers. Although it provides support for the VMware APIs, Isilon has had little success in virtualized environments. EMC has a number of overlapping products (such as VNX, Data Domain, Atmos and Centera), each of which internally competes with Isilon, depending on the use case under consideration. Competition from NetApp has stepped up considerably during the past couple of years, due to a growing focus on Clustered Data Ontap.

Who Should Care: Isilon will continue to be of interest to organizations that have a large amount of unstructured data and are looking for scalable (in both capacity and performance), yet easy-to-

manage, storage with fast sequential throughput. Popular use cases include high-performance workloads, such as genomic sequencing, seismic exploration, financial modeling, rich media broadcast/postproduction, Web analytics, large home directories and active archiving.

Avere

Pittsburgh, Pennsylvania (www.averesystems.com)

Analysis by Pushan Rinnen and Gene Ruth

Profiled in "Cool Vendors in Storage Technologies, 2010"

Why Cool Then: Avere was a startup that launched its high-performance NAS engines in October 2009. Most of the Avere executives worked at Spinnaker Networks, a company acquired by NetApp in 2003 to develop a scale-out NAS solution. Avere offers NAS caching appliances to accelerate file system reads and writes, and to address storage tiering in a more intelligent and cost-effective way.

Unlike other NAS caching solutions in the market today, which address only read I/O performance and are high cost, the Avere FXT appliances address both read and write performance and are offered at competitive prices. Avere developed its technology by examining many workload patterns and comparing them to various storage media characteristics to derive algorithms that automatically allocate the most active data onto the appropriate storage cache media. The appliances embed SSDs, DRAM, nonvolatile RAM (NVRAM) and SAS drives to support all active cache data. All small read and write I/Os are directed to RAM and NVRAM; large random reads are directed to SSDs; and large random writes and large sequential reads/writes are directed to a combination of SAS drives and RAM or NVRAM, where RAM and NVRAM hide access latency. These scalable caching appliances support all third-party NAS systems via NFS for persistent data. Avere also has an integrated, policy-based migration engine that, when the cache is full, migrates data to any third-party NAS system with low-cost Serial Advanced Technology Attachment (SATA) drives. Unlike traditional NAS systems, which scale performance by adding costly FC drives, the Avere appliances scale performance by adding nodes, while storage capacity is scaled separately through low-cost, third-party NAS systems or any existing NAS.

Where Are They Now: Avere has grown into a near-110-employee company, having raised three rounds of funding for a total of \$52 million. The last round was \$20 million in July 2012. Since being profiled in the 2010 Gartner Cool Vendors report, Avere has seen solid adoption in a few vertical industries, especially media/entertainment and high-tech companies that require leading-edge NAS performance. While many customers leverage SATA-based NAS systems from NetApp and EMC as the back end of Avere's appliances, some customers also use open-source ZFS-based NAS systems. Customers of the vendor's latest-generation (version 4) appliances can use object storage in the public or private cloud as the back end.

In 2012, Avere launched its third-generation appliances with two key functions: FlashMove and FlashMirror. FlashMove allows user-transparent data movement and migration between back-end heterogeneous NAS devices, and FlashMirror offers heterogeneous replication. With the latest refresh in March 2014, the heterogeneous back-end storage extends beyond NAS to object storage, such as Amazon public cloud and Cleversafe. When customers use Amazon S3 and

Amazon Glacier together, Avere can keep the metadata and the first 16KB of data of each file in S3 for fast restore with a preview capability. Its FlashMove technology now offers user-transparent data migration from one cloud provider to another cloud provider, and from the cloud to a local NAS device, potentially mitigating users' concerns about their selected cloud provider's viability.

While making good progress, Avere continues to face competition from NetApp and EMC, which try to upsell their large-capacity NAS systems to higher-performance systems. When using cloud storage as the back end, Avere competes against cloud storage gateway vendors; however, Avere, in the role of a gateway, offers unique features that specifically target enterprise customers, including global namespace, transparent migration services and heterogeneous replication services.

Who Should Care: Avere's appliances have expanded its appeal beyond performance NAS customers to customers that also want to leverage inexpensive cloud storage for large archives while maintaining local fast performance, thus reducing on-premises storage infrastructure.

Recommended Reading

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

"Hype Cycle for Storage Technologies, 2013"

"Best Practices in Storage for Server and Desktop Virtualization"

"How to Calculate the Total Cost of Cloud Storage"

"Critical Capabilities for Scale-Out File System Storage"

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>

© 2014 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](#)."