VENDOR PROFILE

Panzura — Cloud Storage for the Enterprise Datacenter

Ashish Nadkarni

IDC OPINION

Globally distributed (geo-dispersed) file-based and object-based storage platforms form the underpinning of next-gen infrastructure. Increasingly, such data persistence platforms are cloud integrated – allowing customers to add cloud as a tier and enable a global data access, distribution, and collaboration mechanism. Early variants of such file systems (i.e., global namespace solutions) are now almost entirely replaced by cloud-enabled storage platforms (CESPs; i.e., cloud gateways) that offer not just cloud as a tier but also essential file system functions such as POSIX interfaces, REST APIs, Hadoop HDFS interfaces, global file locking, access control lists, deduplication, and data encryption. Panzura was one of the first to offer many of these capabilities and continues to differentiate its solution in the market with hybrid cloud storage that front ends public or private clouds with a global file system and a storage controller. Further:

- Panzura has taken care not to let its solution be labeled as yet another cloud-based sharing application for common file access and photo sharing or just a cloud gateway that supports backup and archiving. On the front end, Panzura offers a host of essential enterprise features such as distributed file locking and byte-range locking, unified cloud-integrated global namespace, extended access control lists, continuous snapshots, global inline deduplication, automated storage tiering, user-defined data pinning, and military-grade data encryption. On the back end, Panzura has made its solution compatible with most major public and private cloud providers and standards. With a feature list that rivals those of incumbent on-premises solutions, Panzura's solution makes a compelling case to be put on the must-evaluate list for buyers looking for a simple but robust, easy-to-manage, and all-in-one global hybrid storage solution.
- With its sights set squarely on the enterprise, Panzura is in a position to lead the new market segment of global hybrid storage solutions – a segment that IDC expects to grow rapidly in the coming years as buyers aggressively seek to change the ratio of their capex to opex budgets with a fully managed outsourced storage service. CIOs worldwide are looking at ways to avoid increased IT budget and head count by leveraging the cloud for their infrastructure needs. IDC expects Panzura to initially gain traction in specific industries or use cases – such as global software distribution, manufacturing, federal, media and entertainment (M&E), and architecture, engineering, and construction (AEC) firms – that are more heavily focused on global workflow sharing and project collaboration. As buyers understand and get more comfortable with replacing primary storage with cloud storage, they may expand the deployment scenarios for Panzura to mission-critical environments, similar to the adoption curve of virtualization or deduplication.

August 2016, IDC #US41641116

IN THIS VENDOR PROFILE

This IDC Vendor Profile discusses Panzura and its hybrid cloud storage solution — a solution made up of a global file system and storage controllers that are front ending public and private cloud storage. Using this solution as a backdrop, this document examines Panzura's profile and strategy as well as some of the challenges Panzura could encounter as it tries to differentiate itself in a market that is getting increasingly competitive.

SITUATION OVERVIEW

Definition of Cloud-Enabled Storage Platforms

At a very simplistic level, CESPs (also referred to as hybrid cloud gateways, cloud controllers, or cloud-integrated storage solutions) can be delivered as appliances (physical servers), as software, or as virtual appliances (as a virtual machine [VM] running on any virtualized server infrastructure). In most cases, they reside on the customer premises and translate cloud storage interfaces (SOAP or REST based) to standard datacenter storage interfaces such as block-based (iSCSI and Fibre Channel) or file-based (NFS and CIFS) interfaces. CESPs can be viewed as an on-ramp or a bridge from a datacenter to the cloud. They function as in-band storage devices similar to traditional on-premises disk storage systems. The biggest exception is that unlike traditional disk systems that use physical disks, CESPs utilize the cloud as the persistent storage layer.

CESPs enable companies to integrate cloud storage into applications and workflows without moving the applications or on-premises systems into the cloud. To make this process seamless, CESP vendors have incorporated enterprise features into their solutions. Almost all of the modern CESPs now feature functions like local caching that leverage flash/SSDs, hot data pinning, encryption, deduplication, compression, smart provisioning, and built-in data protection schemes like snapshots and replicas. Many CESPs even present the concept of a global namespace and distributed file systems (with file locking) that enable a series of gateways to scale out across datacenters and geographies, including remote/branch offices. This makes them suitable in environments that require globally dispersed groups of users to collaborate on projects that leverage the same data sets.

Given the advanced set of functions included with CESPs, referring to them as "cloud gateways" does a huge disservice to their value proposition. However, the industry at large has standardized the term *cloud-enabled storage* to some degree, although arguably, it is too simplistic in that description.

Leading Use Cases for CESPs

Irrespective of how these CESPs get packaged, delivered, and licensed by the vendor (as software, virtual appliances, and physical appliances), they all address the fundamental ability to bridge on-premises and cloud storage environments. Doing so allows businesses to leverage such solutions in diverse use cases, including:

- **Capacity overflow:** CESPs serve as a bridge between on-premises and public cloud storage without making a wholesale move. Existing on-premises NAS or SAN storage can be extended by relying on overflow capacity at a cloud storage provider such as Amazon S3 or alternatives. This could be useful to mitigate on-premises capacity growth or handle NAS sprawl and consolidation.
- **Archive/long-term retention:** Most organizations face the need to retain data for long periods of time for regulatory, legal, or general business/data analytics reasons. Moreover, a significant

amount of today's primary storage houses data that is inactive or obsolete, which IDC refers to as "information rot." Cloud storage can serve as a parking lot for older data, which must be retained, either to adhere to regulatory/compliance policies or because a firm is not confident it can delete the data.

- **Backup and disaster recovery (DR):** CESPs can serve as ideal targets for "offsite data" backups – they ensure data is sent to the cloud in a secure and reliable manner. By using CESPs in conjunction with data protection applications, businesses can minimize or even completely eliminate tape infrastructure as well as tape-handling and tape-collection services. The gateway can serve as a local backup target that also doubles up for local recovery. (Within the backup market, this is known as hybrid cloud backup.) Data on the gateway can be replicated asynchronously, protected via point-in-time snapshots or clones, or even transmitted as backup images to cloud storage.
- **Application-specific use cases:** Vendors like Panzura are up-leveling their value proposition to include application-specific use cases like Microsoft collaboration and productivity applications and CAD/CAM applications for architecture, engineering, and construction firms. With unique features such as global/distributed file locking, CESPs claim to have solved the notorious "file open" problem that has plagued firms with globally distributed teams that need to collaborate using the same sets of files.

Packaging and Licensing Flexibility

Vendors of CESPs rely on different packing or delivery mechanisms to cater to a diverse set of use cases or markets that they're pursuing. These include physical appliances, virtual appliances, software, and embedded controllers. In many cases, such decisions are purely driven by the need to closely resemble the delivery and support mechanisms relied upon by traditional storage vendors (given that both groups of vendors are essentially going after the same buyer demographic). Delivery mechanisms commonly deployed are:

- **Hardware appliances:** Appliances incorporate software logic on a storage controller – which in most cases is an x86-based server. The server resides in-band and on-premises between the application and the off-premise cloud storage. Akin to a router on one hand, its "WAN" interface is optimized for TCP/IP-based SOAP or REST communications to support cloud services. On the other hand, it's "LAN" interface is optimized for providing IP-based file or block connectivity to applications that reside in the datacenter. Most cloud storage vendors support this delivery model.
- **Virtual appliances:** Virtual machines that run the CESP software run on the business' on-premises virtual infrastructure. In some cases, they can also reside on purpose-built appliances that support workload-specific virtual machines. The use of a virtual appliance will typically reduce the cost and time of implementation and configuration and is ideal for smaller, lower-performance environments (i.e., remote offices).
- **Software only:** Businesses can choose to run the CESP software on servers that run other applications or workloads. While this packaging mode is uncommon and not preferred by many vendors, they nevertheless offer CESP software for bulk installs.
- **Embedded:** Through partnerships, CESP vendors may choose to embed their code into another appliance or stack. The most common implementation of this model is in personal and entry-level storage (PELS) systems that are sold either as standalone units to "prosumers" or as part of a larger deal in certain verticals like FSI.

Vendors of CESPs have also adopted innovative licensing models to make their solutions financially attractive from an opex perspective – on similar lines as the cloud storage that they enable. Licensing models include (but are not limited to) unlimited licensed capacity for cloud storage, service-based model that is based on managed terabytes, and amortized appliance costs over the life of the cloud storage service.

Company Overview

Panzura was founded in 2008 by a team of engineers with experience in developing enterprise infrastructure technology. The technical team has roots in Oracle-Sun; NetApp; Data Domain; Riverbed; Aruba, a Hewlett Packard Enterprise company (formerly Aruba Networks); and Veritas, among others. Panzura's founders realized that traditional storage technologies could not meet the needs of a global enterprise that requires its globally distributed workforce to collaborate and share workloads at LAN-like speeds while keeping networking and storage costs at a minimum. Panzura's vision is therefore to "transform the traditional enterprise storage model to the cloud and unlock the power of data."

Panzura's flagship hybrid cloud storage (HCS) solution is an internet-based, interconnected, unified global file system based on a purpose-built OS and storage controllers. With the Hybrid Cloud Storage solution, persistent data can be centrally located for ease of management and unlimited cloud storage capacity, and cost control yet remains easily accessible, regardless of physical location.

By deploying Panzura, buyers can seek to:

- Offload cumbersome storage and services management to the cloud, improving productivity and increasing focus on high-value IT initiatives
- Minimize storage islands at various branch offices by deploying a unified solution that centralizes data yet allows users to collaborate and share data as if they were local
- Lower their capacity and data protection costs with a solution that has built-in data protection and active archiving capabilities
- Simplify system management globally by having a single solution for project collaboration and workflow

Company Strategy

Panzura has thus far ensured that its solution remains well differentiated in an increasingly crowded market. It has sought to make its solution, like standard NAS, attractive to buyers of all sizes and suitable for a wide variety of workloads, including desktop virtualization and a wide variety of verticals, such as global software distribution; architecture, engineering, and construction; and manufacturing. Clearly, Panzura is positioning its solution where it can showcase its strength: the ability to centralize storage using public or private clouds yet globally distribute data and make it accessible at local speeds. However, Panzura is positioning its solution directly against incumbent storage technologies with a focus on creating an all-in-one solution with enterprise features such as:

- A global file system and a unified global namespace connecting all sites worldwide and accessible via NFS and CIFS locally, ensuring that clients can communicate using industry-standard protocols and can get a location-agnostic universal view of their workspaces regardless of location
- Distributed file locking and byte-range locking for data integrity, which is an especially important feature in global workflows and project collaboration environments yet avoided by most cloud solutions because of the engineering difficulty in getting it right

- Globally coherent and distributed storage optimization capabilities such as inline deduplication that is unique in the industry, auto-tiering, and WAN optimization
- Data protection and security, including continuous snapshots and military-grade AES-256-CBC encryption, ensuring that data can be protected and recovered regardless of its location, with the encryption key managed and controlled by IT
- SSD-based data caching and policy-based capabilities including pinning to provide LAN-speed access to data
- Hybrid cloud NAS with support for public cloud providers, private cloud providers, and frameworks like Amazon, Microsoft Azure, Google, IBM (Cleversafe), EMC, and OpenStack

Panzura is highlighting the viability of its solution by demonstrating how it works well in specific geo-dispersed environments that buyers find challenging to deploy on traditional storage solutions. For example:

- Active archiving of unstructured data for specific verticals such as healthcare, energy exploration, media/entertainment, and financial services
- Seamless cloud backup of unstructured and structured data enabled by integration with major native tools and backup solutions like IBM Spectrum Protect (formerly Tivoli Storage Manager), Veritas NetBackup, and CommVault
- Ability — as a cloud-enabled NAS platform of choice — to integrate all compute workloads of global distributed environments, providing cloud-based data protection, DR, and instant recovery capabilities, even for virtualized environments

Panzura has seen robust success with major worldwide organizations including Electronic Arts, U.S. Department of Justice, National Instruments, DreamWorks, and Comcast. Panzura controllers are available in physical or virtual instances, with the vMotion support for virtual instances to help lower both the management and the storage costs associated with typical VMware server and desktop environments. Panzura claims that an additional benefit for VDI environments is the ability to create a single global profile that "roams" with the user, providing localized pinning as the user hops from office to office.

FUTURE OUTLOOK

IDC believes that cloud as a tier will continue to make a big influence on how buyers assess storage solutions in terms of systems' performance, operational efficiency, and data intelligence. Companies like Panzura will be at the forefront of this change. They will accelerate the adoption of cloud as a tier by eliminating nagging issues such as data security, resiliency, and data localization and pinning. This increased demand will have a floodgate effect on cloud providers and hybrid cloud NAS solutions like Panzura's.

Another development in this field is an increasing push toward adoption of newer file and object-aware standards. While NFS and CIFS will continue to dominate, and thus the adoption of NFSv4 and SMB3.0 will be a progressive update for file-based solutions, the availability of newer interfaces like CDMI, HTTP/REST, and WebDAV will push suppliers into newer territory. Some already have a jump-start by their object-based legacy, while others with pure file-only interfaces may have somewhat of an uphill climb. Thankfully, the sluggish economy coupled with slower market adoption of these standards means that these suppliers have some time to get into the "next gen" interfaces game.

Finally, from a competitive standpoint, while the hybrid cloud storage market is predominantly made up of start-ups, there is no doubt that the incumbent suppliers are watching this market closely. Given the

history of the storage market, consolidation is one of the strongest possibilities, as is being observed of late in the all-flash array market. The biggest question is whether the acquirers will be cloud providers or incumbent storage suppliers.

ESSENTIAL GUIDANCE

Buyers that are considering investing in CESP should examine the strategic role it will play in their IT infrastructure:

- **Functionality:** CESPs vary greatly in the functionality they offer – and are highly workload dependent. Gateways optimized for block workloads are not well suited for file workloads and vice versa. Similarly, the breadth and level of support for different types of clouds varies greatly. Buyers should also examine the performance impact of having primary applications access data on or via the CESP – technologies like SSD/flash-based caching, pinning of data sets, global/distributed file locking, and data mobility are crucial decision points. Buyers should also examine the federated multicloud data management capabilities as well as SLA-monitoring and enforcement capabilities.
- **Investment protection:** Given that CESP providers remain vulnerable, buyers should generally avoid platform or cloud lock-in. They can do so by ensuring that a copy of their data is stored in a different cloud (or location) and that it can be accessed without the CESP being in the data path. Buyers should also examine the impact this investment may have on existing datacenter storage technologies like scale-out file-, block-, and object-based storage.

Advice for Panzura

The success of Panzura's solution will be largely driven by the traction it gets in the specific market segments and verticals struggling with data deluge. While Panzura's feature list is appealing to almost all market segments, some conservative enterprise buyers that may be a bit squeamish about putting their data in a public cloud may require some convincing. Similarly, Panzura may find it easier to gain traction in certain verticals for which geographically dispersed collaboration, data management, and active archiving/protection for collaborative data sets are must-haves.

Panzura may also have to do a balancing act between competitors and partners. Cloud platform providers and cloud storage providers may also choose to enter the market with global file system and namespace solutions. But this is no trivial exercise. Panzura's technological lead and early moves in this direction should help the company ward off some of these potential threats, but there is always the risk that newer suppliers will enter the market and try to emulate Panzura's technology and go-to-market strategy. Panzura should continue to differentiate based on its core strengths while trying to appeal to a wider buyer base.

LEARN MORE

Related Research

- *A Survival Guide: The Data Protection and Recovery Software Market Transition in the Third Storage Epoch* (IDC #US41605916, July 2016)
- *DATAGRES - Enabling the Transformation to IDC's 3rd Platform Computing* (IDC #US41593316, July 2016)
- *IDC's Worldwide Storage Software Taxonomy, 2016* (IDC #US41593216, July 2016)

- *Datos IO: Data Protection for Next-Gen and Cloud Databases* (IDC #US41584516, July 2016)
- *Cavium to Acquire QLogic for Approximately $1 Billion Enterprise Value* (IDC #US41536116, June 2016)
- *IDC's Container Infrastructure Software Ecosystem Taxonomy and Overview, 2016* (IDC #US41401716, June 2016)
- *IDC's Worldwide Computing Platforms Taxonomy, 2016* (IDC #US40861916, April 2016)
- *Open Compute Project - Full Steam Ahead* (IDC #US41138316, April 2016)
- *Evolving System Architecture: Disaggregated and Composable Infrastructure* (IDC #US40862016, March 2016)
- *Cleversafe Object Storage-Based Cloud Storage Services Enable "Exascale" IBM Cloud Hybrid Storage Architectures* (IDC #lcUS41067816, March 2016)
- *New Technologies Enable New Workloads* (IDC #DR2016_T6_AN, March 2016)
- *Software-Defined Infrastructure in a Composable World* (IDC #DR2016_T6_AG, March 2016)
- *Building an ARM: How Demand for Computing Platform Choice Generated a New Computing Ecosystem* (IDC #US41031016, March 2016)
- *Impact of Next-Gen Applications on IT Infrastructure* (IDC #US41031116, February 2016)
- *Talena - Always on Big Data* (IDC #US41007916, February 2016)
- *IDC's Worldwide Storage for Big Data and Business Analytics Taxonomy, 2016* (IDC #US40997216, February 2016)
- *IDC's Worldwide Server-Based Storage Taxonomy, 2016* (IDC #US40995216, February 2016)
- *Polyglot Persistence: Insights on NoSQL Adoption and the Resulting Impact on Infrastructure* (IDC #US40997116, February 2016)
- *IDC's Worldwide Software-Defined Storage Taxonomy, 2016* (IDC #US40951716, February 2016)
- *Worldwide Software-Defined Infrastructure, 2014-2019: Forecast Report* (IDC #US40903016, January 2016)
- *Worldwide Software-Defined Storage, 2014-2019: Forecast Report* (IDC #US40903816, January 2016)
- *Worldwide Disk-Based Data Protection and Recovery Forecast, 2015-2019*: Blame the Cloud? (IDC #US40662415, December 2015)
- *IDC MarketScape: Worldwide All-Flash Array 2015-2016 Vendor Assessment* (IDC #US40721815, December 2015)
- *Worldwide Storage for Public and Private Cloud Forecast, 2015-2019* (IDC #US40702215, December 2015)
- *New IDC Forecast Sees Worldwide Big Data Technology and Services Market Growing to $48.6 Billion in 2019, Driven by Wide Adoption Across Industries* (IDC #prUS40560115, November 2015)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com