

## REPORT REPRINT

# Primary Data seeks to crush data management challenges

HENRY BALTAZAR

15 MAR 2016

The company's management team helped revolutionize the use of flash in servers at Fusion-io, and is now looking to advance the world of storage management and optimization through the launch of DataSphere.

---

THIS REPORT, LICENSED EXCLUSIVELY TO PRIMARY DATA, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR REPOSTED, IN WHOLE OR IN PART, BY THE RECIPIENT, WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2016 451 Research, LLC | [WWW.451RESEARCH.COM](http://WWW.451RESEARCH.COM)

Primary Data's management team helped revolutionize the use of flash in servers at Fusion-io, and is now looking to advance the world of storage management and optimization with the launch of DataSphere. The company unveiled its DataSphere data virtualization platform at VMworld 2015; it is currently deployed at a few early adopter accounts. Primary Data is looking to advance into production deployments in 2016, and could raise an additional funding round to accelerate its progress.

---

## THE 451 TAKE

With the emergence of flash and object cloud storage services creating new high-performance and cost-efficient data retention tiers, respectively, these new innovations make a stronger case for the use of data virtualization products such as Primary Data's DataSphere. Although the storage industry has tried on multiple occasions to launch data management and migration initiatives, the complexity of these offerings has outweighed the potential benefits. Primary Data believes the rampant data growth and complexity that enterprises are struggling with will create an opportunity for next-generation data virtualization and management tools. Like previous storage optimization startups, Primary Data will have to build up its base of reference customers, as well as market awareness for DataSphere, to convince enterprises to carve out budget for its data management and optimization platform.

---

## CONTEXT

Los Altos, California-based Primary Data has taken in \$63m in venture funding to date, and will likely add to its haul in 2016. Its last funding round was in November 2014 to the tune of \$50m; it was led by Accel Partners, and included Battery Ventures, Pelion Venture Partners, Lightspeed Venture Partners and Wing Capital Group. The vendor says it has a substantial amount of capital left in its coffers, and is seeking an up round. Primary Data's headcount has hovered near 80 for the last year and half, but will likely increase this year as it ramps up its sales to production customers.

The vendor continues to be led by Fusion-io veterans David Flynn, Rick White and Lance Smith, with Apple co-founder Steve Wozniak holding the role of chief scientist. Primary Data unveiled its DataSphere platform at VMworld 2015, and has its technology deployed with a small group of early adopters, with named customers coming later this year.

## STRATEGY

Primary Data claims the key value proposition for DataSphere's data virtualization is the ability to give applications and clients the right data – at the right place and at the right time – while using customers' existing storage. The vendor is targeting a few key vertical markets with its early adopter trials, including telco, software/technology vendors, travel, managed service providers, financial services, and media and entertainment.

Primary Data has a handful of partnerships, including an arrangement with Dell, which is an investor, and is integrating hardware for customers that would prefer to purchase DataSphere in an appliance format. It also has a technology partnership with NetApp, and is helping field teams migrate data from legacy systems to Clustered ONTAP. The company will ramp up its reseller channel in 2016, with a focus on attracting Fortune 2000-class companies. Irvine, California-based Trace3 is the vendor's first integration partner.

The company initially planned on creating an all-flash storage system to act as an acceleration tier, but scrapped the plans for this product since it believed that this additional component would scare away potential storage partners. This was a prudent decision, since it will allow Primary Data to work with a wider range of vendors and eliminates the fear that DataSphere requires a costly rip-and-replace deployment, which storage professionals universally hate.

## PRODUCTS

DataSphere was designed to provide a storage-agnostic management and data migration architecture, with the policy engine and data analytics management components running out of band. The DataSphere manager analyzes the capacity and performance utilization of workloads and initiates data migration operations by sending commands over an out-of-band control channel when the need arises. Primary Data's DSX software runs on each managed host system, and is responsible for migrating data across NAS and SAN storage silos; the hosts use a 10Gbps Ethernet connection to move data horizontally. The vendor claims its DSX data movers have resource consumption intelligence to make sure that application performance is not severely impacted by the in-place data migrations, and its analytics and monitoring tools ensure data is not migrated needlessly, to avoid thrashing.

DataSphere's data virtualization allows organizations to set up a catalog of service levels tuned to match the performance (IOPS, bandwidth and latency) and protection (durability, availability, recoverability and security) needs of applications. The first iteration of the product only supports VMware hosts today, although there are plans for future expansion to Hyper-V and KVM.

Primary Data has not disclosed the pricing for DataSphere, although it has said that licensing will be an annual subscription, with a DataSphere license for each cluster manager and a DSX license for each host that participates. DataSphere will not have any capacity-based licensing, which goes against the grain compared with older storage and data management products.

## COMPETITION

The storage management and optimization space continues to be dominated by storage array players such as EMC, IBM, NetApp, HDS and Hewlett Packard Enterprise, which typically sell their management tools to array customers. The capabilities that Primary Data has brought to market have some similarities to software-defined storage management stacks such as EMC's ViPR and IBM's Spectrum Control. Primary Data claims its differentiation is the ability to work with any storage platform, while established competitors have focused on their own arrays. It also notes that most early software-defined storage deployments are more focused on monitoring and provisioning, and not on performance optimization.

In the management space, there is a growing crop of players that are also looking to help customers that are struggling with storage management and optimization, including ATS Group's Galileo, Datagres, Load DynamiX and Virtual Instruments. There are a couple of vendors and specialists that have data migration tools, including Data Dynamics and Komprise, as well as older storage virtualization players such as DataCore Software and FalconStor Software. In the software-defined-storage space, there are several startups looking to attack similar storage cost-reduction and -efficiency issues, such as Formation Data Systems, Hedvig and Nexenta. The differentiation for Primary Data will be its data migration capabilities, which not only cover on-premises storage assets, but also provide a path to cloud storage services.

## SWOT ANALYSIS

### STRENGTHS

Primary Data has a proven management team that has experience establishing new market categories, such as PCIe flash storage, at previous ventures. The vendor was able to raise a relatively large amount of funding initially.

### WEAKNESSES

The vendor still needs to establish a track record in production environments, and must flesh out its go-to-market capabilities. Data Sphere is limited to VMware environments in its initial release.

### OPPORTUNITIES

The burden of storage management continues to get heavier, and tools like Primary Data's DataSphere can help organizations boost the efficiency of their storage infrastructure.

### THREATS

A growing number of startups and established players are developing storage management and optimization tools to address the same set of pain points. Primary Data will need to ramp up marketing efforts to increase awareness and differentiation.