

## REPORT REPRINT

# Portworx is driving cloud-native storage and data management for Kubernetes

**NOVEMBER 07 2019**

**By William Fellows**

The company's view is that in order to make Kubernetes a first-class citizen in the enterprise tool chain - to move from PoC into production - Kubernetes will need to move data faster and make it safe. Portworx is growing at a clip with an impressive set of reference customers and oversubscribed investor interest.

---

THIS REPORT, LICENSED TO PORTWORX, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT 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 RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



## Introduction

The cloud-native adoption barrier is low enough that users can get a lot of value from using containers with the low-hanging workload and application fruit. However, the next set of applications is going to have more demands around data security, portability, encryption, backup and storage – these are requirements that Kubernetes alone can't meet. Portworx's view is that in order to make Kubernetes a first-class citizen in the enterprise tool chain – to move from proof of concept (PoC) into production – it will need to move data faster and make it safe.

### 451 TAKE

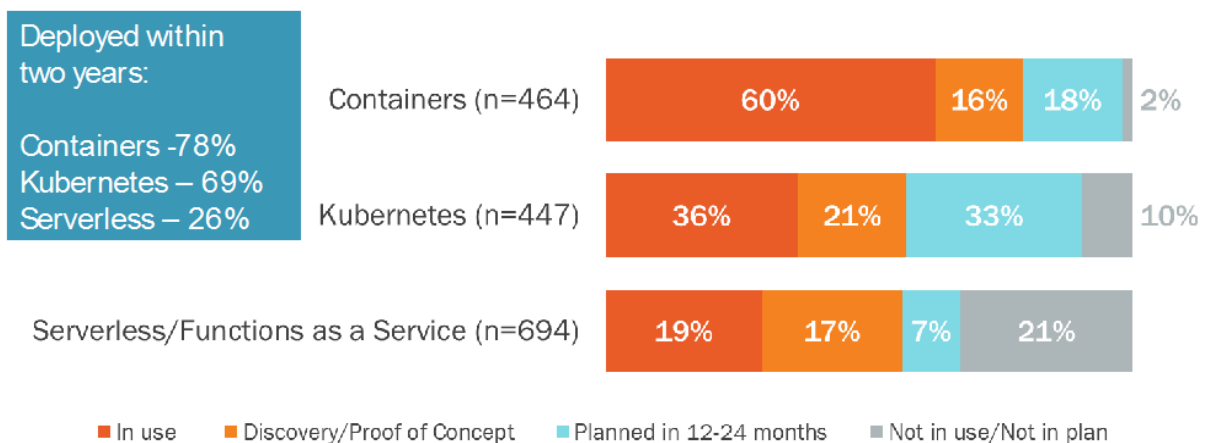
As cloud-native moves from the project and PoC stages into production and to the overall IT strategy, enterprise-class storage and data orchestration will be required in addition to Kubernetes container orchestration (see figure below). The business problem Portworx is aimed at is data management for applications running on Kubernetes and enabling reliable deployment of containerized enterprise applications to persistent storage. The company is growing at a clip with an impressive set of reference customers and oversubscribed investor interest.

## Context

Portworx says its customers are typically building an application platform (either SaaS or PaaS) with several design points: self-service for developers, automated, infrastructure-agnostic SLAs, low-touch operations and cost-optimized. However, Kubernetes can't run business-critical applications alone for the reasons above – it wasn't designed to do those things. This means customers' legacy platforms live on (and expand). Portworx finds that customers typically try to use storage array connectors or existing storage offerings, let the database handle data in containers, or use block storage provided by the cloud. Connectors are built for VMs, have an operational overhead, use NFS and are usually a single availability zone. Databases are expensive and don't do backup, security and upgrades. They also increase compute costs by overprovisioning compute, Portworx finds. Block storage is also designed for VMs, is locked to a cloud provider and doesn't have data management capabilities.

### Adoption of Cloud-Native Technologies

Source: 451 Research's Voice of the Enterprise: DevOps Q1 2019 and Digital Pulse, Budgets & Outlook 2019



### Product

This is where the Portworx Enterprise Data Platform comes in. It's a container-only offering (not for VMs) and is used for internal PaaS (T-Mobile); SaaS apps (ESRI); multi-cloud Kubernetes deployments (Iceland Bank); IoT and real-time processing (NIO); and DR for Kubernetes (a NYC bank). The platform comprises PX-Store, which provides reliability and performance for container volumes; intelligent scheduling engine; HA and data replication; local snapshots with automated restores; storage performance tiering; shared volumes (ReadWriteMany); and automatic failure domain detection. PX-Central provides monitoring and management including control center of Portworx Enterprise; cluster data and metric aggregation; preconfigured with dashboards; management console for storage operations; and built-in alerts, and it integrates with Sysdig and Datadog.

PX-Security offers data security for the container platform; bring-your-own-key data encryption; integration with leading secret stores; granular data volume access control; and integration with corporate identity services. PX-Data Management delivers multi-cloud and multi-cluster data and app migration; application-consistent snapshots; off-site snapshots to cloud and object storage and five-minute RPO. It captures the whole app for migration and recovery (such as blue-green deployment), as well as for lift/shift, maintenance, clean installs and augmenting capacity. PX-DR provides multi-site synchronous replication; RPO-zero failover in the same metro area; multi-site async replication in the same WAN and continuous backups across global datacenters. It enables customers to move their cloud-native stack (Kubernetes, scheduling, execution runtime (e.g., Docker), CSI, CNI and networking across different clouds (compute, storage, network) and to bare metal as either migration or bursting.

### Business model

In March, Portworx closed a \$27m series C funding round co-led by Sapphire Ventures and the ventures arm of Mubadala Investment Company, with support from existing investors Mayfield Fund and GE Ventures, and new financing from Cisco Investments, HPE and NetApp. Total funding is \$55.5m. We estimate the 65-person company has about \$8m in revenue. The company was founded in 2015 by veterans of Google, Microsoft and Nimble Storage, and by the same founders as Ocarina Networks, which was acquired by Dell in 2010. It says that every initial \$1 of customer spending becomes \$5 within two years. It claims 13 deals that started at over \$100,000 and that 30% of its orders are reorders. It now has 100+ customers (45 in the Fortune 2000) including GE, Lufthansa, T-Mobile, Adobe and Comcast.

### Competition

Portworx's competition includes Oracle's Autonomous Data Platform, DIY approaches using Ceph and vendors moving to cloud-native from the VM space. Other vendors in the cloud-native storage space include Diamanti and Robin Systems. Most of the major storage and systems vendors now offer container-based storage.

SWOT Analysis

**STRENGTHS**

NetApp, BMC, VMware vSAN and other storage environments are all enterprise-class, but they are not designed from a container point of view and do not speak Kubernetes natively. If Kubernetes makes applications agile, Portworx is aiming to do the same for data in cloud-native environments.

**WEAKNESSES**

The cloud-native market is thrashing with players, and the noise level is high. Portworx will need to stay ahead of the cloud-native pack, as well as the legacy and incumbent data management companies as they hone in on this opportunity because the container will be the atomic unit of the current re-platforming cycle.

**OPPORTUNITIES**

At its current course and speed, Portworx looks as if it will go quickly from being an interesting player to a scale player of significance.

**THREATS**

Portworx will need to have the major systems, services and storage vendors as partners rather than face them as competitors. The question is whether there is an opportunity for an independent, growing, long-term, sustainable and profitable company, or will this class of cloud-native data and storage management likely be consolidated into broader management offerings.