REPORT REPRINT

VM data protection veteran Zerto sets its sights on Kubernetes-based stateful apps

OCTOBER 20 2020

By Liam Rogers

Zerto is bringing its granular journal-based data protection to the world of Kubernetes via the yet-to-be-released Zerto for Kubernetes platform, which takes familiar Zerto data protection capabilities and wields them in support of stateful container applications.

451 Research®
Now a Part of
S&P Global Market Intelligence

 | WWW.451RESEARCH.COM

Introduction

Zerto has built a reputation for highly granular disaster recovery (DR), and over the past couple of years the vendor has delved further into backup and long-term data retention as it has reinforced its core platform. Now it is looking to build alongside that platform as it develops Zerto for Kubernetes (also known as Z4K), which takes familiar Zerto data protection capabilities and wields them in support of stateful container applications.

451 TAKE

Zerto has seen success with its core platform, and now the vendor is porting the same principles that make that platform effective to the world of containers, to provide the same style of granular application-consistent data protection for Kubernetes as it does for VMs. By sticking with the journal-based data protection of its core platform, Zerto hopes to gain an edge over vendors relying on a less continuous snapshot-based approach. The addition of Z4K to its portfolio addresses an emerging market and enables it to target the overlap of containers and VMs, to serve customers that have to manage the backup and DR needs of both technologies as they are intermingled. Zerto is also targeting the protection of stateless applications; although there might not be persistent volumes to recover, the goal is to provide a faster way to restore apps than rerunning from a git repo. As a pure-play data management vendor, there are ample opportunities for partnerships with on-premises and public cloud primary storage providers, although there will be an element of 'co-opetition' with some as other vendors turn increasing attention to Kubernetes services.

Context

Zerto was founded in 2009 and is dual-headquartered in Boston and Israel. The company boasts over 8,000 customers and has 700 employees. Zerto has taken in $183m in funding to date. Its most recent round was a series F of $33m in June. The company intends to expand its continuous data protection capabilities around SaaS, PaaS and containers. Z4K accomplishes this on the Kubernetes front and is aimed at customers that are further along with their container deployments and realizing the challenges of protecting persistent data volumes for Kubernetes. In our Voice of the Enterprise: Storage, Data Management & Disaster Recovery 2020 study, 43% of organizations with containers in use indicated that they rely on legacy or preexisting data protection tools as their primary data protection strategy for containerized applications and associated data volumes. Although such tooling is still important for protecting VMs, dynamic cloud-native applications require a different approach that necessitates new tooling.

Strategy

Z4k is currently in alpha, and Zerto anticipates a public beta before the end of the year, with general availability early in 2021. The alpha has been in conjunction with select partners offering major Kubernetes services, including Red Hat, AWS, Microsoft Azure, GCP, IBM and VMware. There are currently 10 alpha users made up of existing Zerto customers, including tech partners and managed service providers. Zerto has a significant MSP partner base and says that interest in Z4K from these partners goes hand in hand with more of these providers offering Kubernetes services and wanting data protection capabilities they can leverage internally to support those services.

The audience for containers and Kubernetes includes a variety of roles that would historically be outside the bounds of the typical Zerto buyer. This challenge is not unique to Zerto, and all data protection companies playing in the Kubernetes space are faced with having to learn to speak to new personas that are potentially unfamiliar with routine data protection. Additionally, management functionality must cater to different audiences and their skillsets, whether the preferred method is a GUI, API or command line interface (CLI) . Zerto will have to ensure it makes Z4K accessible to developers both as a product and in the function of the software.

Product

Z4K offers continuous data replication for backup, DR and mobility for Kubernetes-based applications that use persistent volumes to store data, whether they are running on-premises or in the cloud. Z4K takes the same journal-based workflow for point-in-time restores that Zerto's flagship product is known for and applies it to Kubernetes. While the core IP around the Zerto journal is intact, its form factor has been altered. The journal-based replication engine runs as a DaemonSet at the worker-node level within a cluster while a Zerto Manager pod runs on any node and communicates with a cloud-based Z4K manager for global availability. The cloud-based manager can then be engaged via API or CLI. Although Zerto is leveraging some of its core IP here, this approach involves rewriting the management components so that they can communicate with the Kubernetes API.

The product is currently installed via the Kubernetes CLI through YAML files, but the vendor intends to have a Helm Charts and a Kubernetes Operator available at GA to ease installation and management. Z4K leverages the existing Kubernetes Storage Classes, which are already baked into the Kubernetes CSI framework as a way to simplify setting up volumes used for recovery and journaling. Once a persistent volume claim is submitted, the platform can begin journaling any chosen application's persistent data into a journal volume, where it will capture the entire state of that Kubernetes application along with all changes. The service enables IO consistency across multiple persistent volumes so that it is crash-consistent across multiple pods. The journal's ability to take frequent checkpoints (recording delta changes as frequently as seconds apart) enables highly granular recovery and presents an alternative to snapshot-based data protection, which can result in gaps in protection from the time between snapshots. Tagging can be done to identify known good checkpoints to make them easier to identify in the checkpoint log, to facilitate swifter recovery. By using a mechanism for tagging checkpoints, a user can gain application consistency and recover or restore to a known good point in time.

The pricing model for the service has yet to be finalized, but it will be offered as a separate product for customers that aren't already leveraging the main Zerto platform. Zerto has also said it is considering offering Z4K as SaaS, wherein Zerto would be deployed locally but management would be done via a central portal.

Competition

Zerto is not the first vendor to tackle backup and DR for stateful apps. Kasten was an early entrant into data protection for Kubernetes-based apps, and the startup was recently acquired by backup heavyweight and Zerto competitor Veeam. Because Veeam has been a longtime rival of Zerto in the VM backup space, we can expect that the combination of Veeam and Kasten will continue to present direct competition. Portworx, which was recently acquired by Pure Storage, has its own Kubernetes backup service that customers can purchase without having to buy into the broader Portworx storage platform. NetApp's Project Astra, announced earlier this year, is a platform for cloud-native app-data lifecycle management. These vendors all cater to backup, DR and application mobility. VMware offers Kubernetes data protection within Tanzu, building on the Velero open source project the vendor supports. Other vendors providing data management capabilities for Kubernetes-based apps include Arrikto, Cohesity, Diamanti, Mayadata, Robin and Trillio.

SWOT Analysis

STRENGTHS

Zerto has a well-established foothold in the data protection market and has a stable of customers, including a significant number or MSPs, to cross-sell into as their cloud-native data protection needs grow.

WEAKNESSES

Zerto has been slower to the Kubernetes space than some of its competitors, so it will have to work to establish the new service and lean into its points of differentiation.

OPPORTUNITIES

As of this writing, the workflow in Z4K is terminal-driven and the vendor will want to implement a simplified GUI-based workflow to go after less operations-skilled audiences.

THREATS

Kubernetes data protection is no longer confined to the realm of startups, and some of the big primary storage vendors that Zerto has typically partnered with are wading into this sector, making for a complicated competitive landscape.