Vendor Profile

Trilio: Cloud-Native Data Services for Kubernetes with a Focus on Application Resiliency

Lucas Mearian

IDC OPINION

Container deployments are forecast to grow at high rates over the next several years, particularly among large enterprises as they make the shift to modern, cloud-native applications running on orchestration engines. IDC finds:

- A Kubernetes cluster may be made up by 100+ nodes. At any given time, Kubernetes will automatically move containers around a cluster to ensure the containerized applications have the compute, memory, and storage resources it needs as workloads evolve and change. Therefore, containers are constantly being killed off, migrated, and spun up based on workload requirements. DevOps teams also often move Kubernetes environments between clusters, both on premises and in public clouds, as containerized applications are created, tested, and brought into production.
- As containerization technology evolved and began supporting traditional applications, the need arose to store the data created by the containerized application. The state of the application, which includes configurations and the session IDs, also needs to be stored so that upon restore it can continue running as previously.
- One of the big challenges around containers and Kubernetes in general is protecting and migrating the persistent storage. An administrator could go and protect the enabling storage system (and associated metadata) on a storage volume; they could also back up the VMs that Kubernetes is running in. But, in both the previous cases, there'd be a lot of stateful data that applications need to run that wouldn't get backed up because it's not persistent.
- A Kubernetes-native application has many components including CSI-attached storage volumes, pods of running containers, and storage resources. Storage resources include persistent volume claims including Kubernetes Secrets (which store and manage passwords), OAuth tokens, and hash keys. All these individual components need to be backed up.

IN THIS VENDOR PROFILE

This IDC Vendor Profile examines Trilio's software for Kubernetes environments that offers backup and recovery of the entire containerized application, including data, metadata, and any other Kubernetes objects associated with the application. This ensures the containerized application is protected and able to be restored from any point in time across multicloud infrastructure.

August 2021, IDC #US48131821

SITUATION OVERVIEW

Company Overview

Founded in 2013, Trilio's TrilioVault Data Protection platform supports Kubernetes, OpenStack, and Red Hat virtualization environments. The TrilioVault for Kubernetes (TVK) software offers cloud-native backup and recovery of an entire containerized application. The backup application is cloud native, meaning like the Kubernetes environments it backs up; it is also containerized for portability and CI/CD enablement.

TrilioVault for Kubernetes backs up a containerized application, as well as its state, metadata, and objects associated with the application and the data created by the application itself. TVK supports applications that are deployed using labels, operators, Helm charts, or namespaces. Backing up the entire application and its state enables TrilioVault to perform asynchronous replication from one Kubernetes cluster to another, or in other words, a point-in-time application recovery.

The company was by founded by its current CTO, Murali Balcha, who previously worked for EMC as a lead development engineer of storage, virtualization, and cloud computing systems. The privately held company is based in Framingham, Massachusetts, and its product is sold globally in the Americas, EMEA, and APAC regions.

Backed by private investors, Trilio has garnered more than $22 million in venture funding to date. Primary investors include .406 Ventures, SKK Ventures, and Plug and Play Tech Center.

The company has about 95 employees and claims to have "hundreds" of customers worldwide, with the majority of them in the telecommunications, defense, automotive, and financial services vertical markets. Predominantly, the software is used in those industries to recover from disasters, migrate workloads, move workloads to new infrastructure, and migrate to new software distributions. The software can be managed via command-line interface or graphical user interface, making it amenable to both DevOps and IT operation teams.

Trilio totes its agentless software's enablement of portability of application backups based on it being a cloud-native application. It captures the state of a containerized application and associated data and allows a user to migrate that to another cluster, either onsite or in the public cloud, where it can be used to restore the previous state as a form of disaster recovery or for test/dev purposes. TrilioVault for Kubernetes supports any Kubernetes distribution, any cloud, any storage, and all Kubernetes application deployment types such as labels, Helm, operators, or namespaces.

For DevOps, the software can restore a point in time of a production application, including its associated data for test and development needs. It can also use a point-in-time copy for identifying performance bottlenecks and troubleshooting data corruption and other issues.

The TrilioVault Management Console discovers applications across various Kubernetes distributions and allows admins to apply data protection policies while managing and monitoring backup and restore plans across hybrid cloud and multicloud environments. The Console allows administrators to define the containerized business applications they want to recover; in other words, the software maps data back to the application instead of a cluster or series of VMs, which would create a more manual recovery process. Trilio's software enables one-button recovery of a single app or a series of apps on a cluster.

In addition to app portability, Trilio is marketing its product as a solution to protect and recover against ransomware attacks based on its point-in-time backups and restores, which can combat data corruption issues or other malicious activity on production data.

Other features as part of the software's capabilities include malware scanning (i.e., using machine learning [ML] algorithms to scan incremental backups and detect if they're increasing in size; if detected, the software will flag that activity and send notifications to the production system). Administrators can then use uncompromised backups to restore business operations. This should occur first in an isolated environment to validate the backup and then again within the production environment once the instance is up and running.

TrilioVault for Kubernetes is deployed as an operator and certified for use with Red Hat OpenShift, IBM Cloud/IBM Kubernetes Engine, VMware Tanzu, HPE Ezmeral, and SUSE Rancher Kubernetes environments. TrilioVault for Kubernetes can be purchased in the Red Hat Marketplace, IBM Cloud Catalog or with IBM Cloud Paks, SUSE Rancher Apps and Marketplace, and VMware Marketplace.

In May, Trilio announced v2.1 of its TrilioVault software, which enabled users to monitor backups by Velero open source software from the TrilioVault Management Console. The new functionality gives users a single view and insights into the status and performance of TrilioVault and Velero backups of Kubernetes resources.

Separately, the company's data protection technology supports OpenStack clouds, and it has been licensed and incorporated into Veritas' NetBackup product to provide backup and recovery, migration, and mobility for OpenStack workloads.

V3Main Technologies

V3Main Technologies, a managed services company based in Houston, Texas, spends much of its time lifting and shifting clients' legacy applications to the cloud and improving cybersecurity and performance for those user's workloads. The company creates custom software and oversees content management systems, ecommerce, websites, and security tools for its customers across various industries, which includes government and public sector organizations.

V3Main uses Kubernetes to containerize client and its own applications to enable microservices-based DevOps model. Its greatest challenge is how to manage its Kubernetes clusters and the storage volumes supporting them in a hybrid cloud environment. Venkat Maddikayala, president of V3Main Technologies, said for the past six months he's been using TrilioVault for Kubernetes for cloud-native backup and restore and for migrations across multiple clouds and multiple Kubernetes clusters.

"Even before I knew about Trilio, when I was working in Kubernetes, my problem was how to take backups of this data. That's where my bottleneck was," Venkat said.

V3Main is currently using all the major hypervisor Kubernetes services, including AWS EKS, Azure AKS, and Google GKE, and is planning to deploy Red Hat OpenShift as well.

TriloVault enables Venkat's team to take a backup from one Kubernetes cluster on AWS where the application is running, for example, and restore it to a different cluster on another cloud service, such as Azure AKS, without making any changes to the application.

"Overall, the concept is to take backup from one cluster, restore it to the target, and then place it on a different cluster," Venkat said. "That's the thing that saves me a lot of time because it takes a lot of time for me to configure a new cluster and on-demand workloads. That's the value."

Pricing

Trilio uses a subscription-based licensing model based on the number of compute nodes, VMs, vCPUs, or clusters in a user's environment.

TrilioVault for Kubernetes list price is $9,000 per year per cluster (a small cluster average is usually around 15 nodes; $1,000 per year per node [volume-based discounting applies]); or $100 per year per vCPU (volume-based discounting applies). Included in the subscription pricing is Trilio's 24 x 7 support and updates. The company also offers special pricing for MSPs and channels.

TrilioVault for Kubernetes is offered as a free trial and free basic edition as well. Both are full-featured versions of the platform.

Company Strategy

Trilio uses a direct sales team leveraging a robust technology alliance partner (TAP) and VAR ecosystem, which the company said helps drive advocacy, solution sales, and ability to transact in local markets.

FUTURE OUTLOOK

The need for persistent storage for containerized stateful applications is quickly growing because container orchestration systems such as Kubernetes help developers build, test, and deploy applications more quickly and with far less effort than traditional operations. IDC expects nearly 2 billion container instances to be installed by 2023, with about 71% of them on non-hyperscaler datacenters. IDC also expects a variety of enterprise workloads such as content and collaboration, business applications, and data management to be deployed on containerized environments.

While the support for enterprise workloads on cloud-native environments is not a new capability, up until recently, much of the innovation has come from start-ups, such as Trilio and its competitors in the container-native marketspace, which include StorageOS, Diamanti, Robin Systems, and Portworx (now a part of Pure Storage). That is changing, however, and traditional IT providers are increasingly adding persistent storage offerings using the CSI standard and data protection services.

As enterprises undertake their DX initiatives, storage and data services associated with containerized environments, microservices, and DevOps will become increasingly important and mainstream. While many container projects today are still in pilot or PoC phases, that is quickly changing with production rollouts. It is therefore critical for container infrastructure providers to remain relevant by providing all the necessary support elements for this increasingly strategic technology market.

ESSENTIAL GUIDANCE

Advice for Trilio

Trilio faces competition from other cloud-native data management and protection start-ups and traditional data protection vendors. Incumbent data management vendors, while offering products that

natively handle Kubernetes environments and the scaling challenges that come along with them, are behind the curve.

Dell-EMC, Red Hat, Hewlett Packard Enterprise (HPE), NetApp, and other vendors are pursuing the container storage market with mixed results, but their development efforts are ongoing. As the shift to cloud accelerates, hyperscalers such as Amazon, Microsoft, and Google will also increasingly offer new tools to simplify the journey to a container-native environment. Currently, start-ups offering persistent storage, backups, and data migration for containerized applications have a leg up on traditional vendors, but as is typically the case with emerging technology markets, that gap will close relatively quickly. Therefore, solutions with a more complete portfolio of service offerings will not only have the potential to attract buyers, but a more comprehensive offering will also bode well for sustainability in the market.

Since the use of containers allows applications to be split into potentially many components, each of which runs in its own container, these environments can increase the number of data objects (i.e., storage volumes) that have to be managed by one to two orders of magnitude (relative to traditional or VM-based environments). Traditional data governance approaches do not scale well into this range and do not offer the granularity needed to manage data effectively in these types of environments.

For example, a traditional Microsoft SQL Server database running in a VM may require two storage volumes (one for the log and one for the data). A microservices-based implementation of that same database may have up to 100+ separate containers collaborating to provide the same application "service." Traditional enterprise storage includes the ability to automatically discover new containers/apps for DevOps. Without knowing where the applications and associated data are, it's impossible to ensure it is protected by the proper data services, and it can be effectively restored.

As containers also exist in a disaggregated architecture — in both on-premises and public cloud environments — there is a growing concern over security, and not knowing where data resides could leave it open to unauthorized access, theft, or manipulation. The location of containers, and the data their applications generate, enables artificial intelligence (AI) and machine learning to make use of that data for analysis. Understanding where data sets reside opens them up to analysis, but in order for that data to be useful, it must first be classified and identified using metadata so that it can be cataloged and made easy to find.

Data discovery is about the ability to find the right data at the right time; this is a key pain point for data engineers and data scientists. Classically, this has been the realm of enterprise search. Most enterprises see AI-enabled search as a key asset for research, analysis, and decision making. Search systems include departmental, enterprise, and task-based search and discovery systems as well as cloud-based and personal information access systems.

Bad actors frequently try to penetrate the backup system either through the administrative console or through the storage media in order to access and delete point-in-time data. Currently, Trilio predominantly provides containerized application backup and data migration capabilities, but the company plans to add additional ransomware protections in its next release (v2.5), which is due out in September 2021. That edition of TrilioVault is expected to include data encryption and immutable backups via an object-locking mechanism on storage media that would prevent backups from being overwritten or deleted. Future editions are expected to help TrilioVault align even more closely with the NIST Cybersecurity Framework, which is aimed at employing existing standards, guidelines, and best practices to reduce risk.

As container production deployments grow, there will be a need for data protection that includes detection and defense against ransomware, as traditional data protection methods do not scale well in containerized environments. Trilio should focus its marketing efforts on what adds differentiation from other competitors, and containerized application security has been shown in IDC surveys to be at the least a big concern among organizations considering or using Kubernetes.

LEARN MORE

Related Research

- *Ionir: Kubernetes Cloud-Native Storage Vendor* (IDC #US47956221, June 2021)
- *StorageOS: Kubernetes Enterprise Storage Vendor Profile* (IDC #US47582821, April 2021)
- *IDC Innovators: Containerized Application Storage Platforms, 2020* (IDC #US46800920, September 2020)
- *Container Infrastructure Software Market Assessment: x86 Containers Forecast, 2018-2023* (IDC #US46185620, April 2020)
- *Why Persistent Storage Matters for Your Containerized Applications* (IDC #US45521719, September 2019)
- *IDC Market Glance: Container Infrastructure Software, 1Q19* (IDC #US44146619, February 2019)

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

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com