

Software Firm Uses Kasten K10 for Kubernetes Backup and Disaster Recovery in a Hybrid Cloud Environment



ENVIRONMENT

- › Google Kubernetes Engine (GKE)
- › Using Google Cloud and on-premises data center
- › Polyglot persistence (NoSQL and relational databases)

CHALLENGES

- › Lack of an application-centric backup and recovery solution for Kubernetes
- › Inability to support hybrid-cloud disaster recovery
- › Business risk arising from data loss or outage

SOLUTION

- › Application-centric data management platform that completely captures a Kubernetes application
- › Simple and easy-to-use software-only solution that spans hybrid and public clouds
- › Policy-driven automation delivering reliable backup operations

RESULTS

- › Unblocked production use of Kubernetes by supporting all compliance requirements
- › Seamless recovery for the entire application in both public and private clouds
- › 70% operator time savings from automation

ABOUT THE CUSTOMER

The customer is a software development firm based in North America with a reputation for innovative design and cutting-edge development. Their projects include security and cryptocurrency solutions and web-based applications with a global footprint.

The DevOps team, that deploys and manages their business critical applications and infrastructure, is also responsible for supporting Kubernetes applications running in Google Cloud including protection against accidental and malicious data loss. For compliance, they also need to replicate and test application backups in a private data center.

LACK OF A DATA PROTECTION SOLUTION FOR KUBERNETES

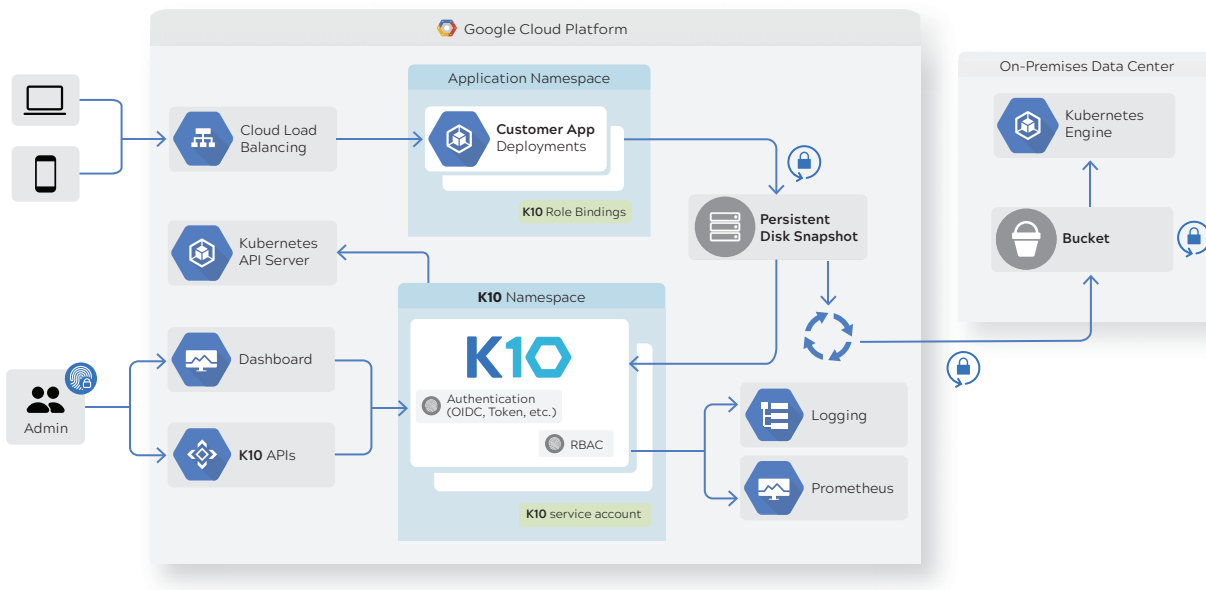
Before discovering Kasten, the operations team ran into two roadblocks while implementing a data protection strategy. First, they needed a solution that captured the complete state of a Kubernetes application. This not only includes data stored on disks but also distinct Kubernetes objects (e.g., ConfigMaps and CustomResources). Second, their desired backup solution was required to operate in Google Cloud and be able to run in and move data between a public cloud and on-prem environment for Disaster Recovery.

Further, time to market was critical. Implementing a data protection strategy could not involve developer or application changes or need manual backup plan modification whenever the application changed (e.g., the addition of a new stateful component).

CLOUD-NATIVE DATA MANAGEMENT WITH KASTEN

In Kasten's K10 platform, the ops team found a solution to both of the above problems. K10, purpose-built for Kubernetes, captures all components of a Kubernetes application including disk snapshots and backups. Further, K10's software-only architecture allows it to run in both Google Cloud and on-prem environments. K10's data portability feature allows any application state captured to be moved seamlessly from a public cloud to an on-prem environment while meeting RPO targets.

Finally, K10's automated application discovery ensures that no developer or application changes are required. The customer's entire applications are auto-discovered and K10 instantly modifies its execution plans to automatically capture any application changes that might occur.



GKE/ON-PREM: HYBRID CLOUD DEPLOYMENT ARCHITECTURE

KEY BENEFITS

- › **Ability to Work with Dynamic Applications Without Developer Overhead:** Without requiring any developer changes, K10 auto-discovers applications, adapts to changes, and dynamically maps policies to the current state of the application.
- › **Easy-to-Deploy Software-Only Platform:** The software-only K10 data management platform can be deployed within minutes to any public cloud or on-prem Kubernetes installation. It also requires no training or professional services.
- › **Policy-Based Backup Workflows:** K10 manages backups at scale through automation and dynamic policies. This avoids the need for custom scripting and allows operations teams to easily create both broad and custom policies for data management compliance.
- › **Improved Time to Market:** K10's powerful workflows, low management overhead, advanced web-based user interface, and centralized monitoring support significantly reduced time spent on data management tasks and managing storage infrastructure.
- › **Infrastructure Independence:** K10 enables disaster recovery workflows across different public or private cloud infrastructure of the entire application stack with seamless data conversion between infrastructure formats when needed.
- › **Full Application Capture:** The ability of K10 to capture the entire application stack by taking a consistent application-to-infrastructure view helps meet critical compliance and restore testing requirements.

Contact Us

289 S. San Antonio Road
Suite 204
Los Altos, CA 94022
United States

+1 (415) 851-1767
info@kasten.io
www.kasten.io
@kastenhq

About Kasten

Kasten is tackling Day 2 data management challenges to help enterprises confidently run applications on Kubernetes. Kasten K10, a data management platform purpose-built for Kubernetes, provides enterprise operations teams an easy-to-use, scalable, and secure system for backup/restore, disaster recovery, and mobility with unparalleled operational simplicity.