



Kubernetes and Cloud Database Backup as a Service

Why CloudCasa for AKS, EKS and GKE

CloudCasa is a powerful and easy to use Kubernetes backup-as-a-service for DevOps and IT Ops teams, that is integrated with Azure Kubernetes Service (AKS), Amazon Elastic Kubernetes Service (EKS), and Google Kubernetes Engine (GKE). Integration with these managed Kubernetes services and your public clouds accounts as well as support all major distributions and managed Kubernetes services enables rapid, cross-cluster, cross-account, and cross-cloud recovery and migration for clusters and their data. CloudCasa is unique in the industry with these integrations that provide hybrid and multi-cloud management and recovery.

"CloudCasa is addressing the key data protection and data migration challenges in multi-cloud and multi-cluster Kubernetes environments," said Enrico Signoretti, Research Product Manager, GigaOm. "With native integration into the Kubernetes services of the top 3 public clouds in addition to Kubernetes itself, CloudCasa helps enable both the multi-cloud and hybrid-cloud journey in this important cloud native ecosystem."

With CloudCasa, you don't need to be a storage or a data protection expert to backup and restore your applications on Kubernetes. CloudCasa does all the hard work of protecting cluster resources and persistent data from human error, security breaches, and service failures, providing the business continuity and compliance that your business requires. It's easy to get started with a few clusters for dev/test, and just as easy for large multi-cluster and hybrid cloud environments. CloudCasa provides application-consistent backups via application hooks with templates, and protection of common adjacent workloads such as Amazon RDS.

In addition to being application and cluster aware, CloudCasa is cloud platform aware. Users can auto-discover clusters on AKS, EKS and GKE, and set global backup policies across a multi-cluster, multi-account and even a multi-cloud environment.

Solution Highlights

As a SaaS offering, CloudCasa offers many benefits vs. do-it-yourself backup software solutions such as Kasten, Trilio and Velero, including:

- Zero infrastructure is required, with a first backup taking less than 10 minutes.
- Best TCO that is supported by a fair, capacity pricing model.

- Backup to user’s own S3 storage with ability to keep all data on premises.
- Best AKS, EKS, and GKE integration for multi-cloud data protection and recovery.
- Cyber-resilient backup service with tamper-proof recovery points in a virtual air-gapped domain, encryption in transit and at rest, and support for private links..
- Available in many marketplaces and catalogs: AWS and Azure partner catalogs, DigitalOcean Marketplace, Red Hat catalog, SUSE Rancher catalog, etc.

Built for the Cloud

Cloud Native: Built with Kubernetes in public cloud, leveraging Kubernetes APIs and the CSI framework to deliver a highly available backup service.

Cloud Scale: Backup service auto-scales to adapt to bursts and growth. Service can also backup to CloudCasa managed storage that never runs out.

Cloud Aware: Integrates with Azure and AWS accounts to auto-discover and auto-protect AKS clusters, Amazon EKS clusters and Amazon RDS databases.

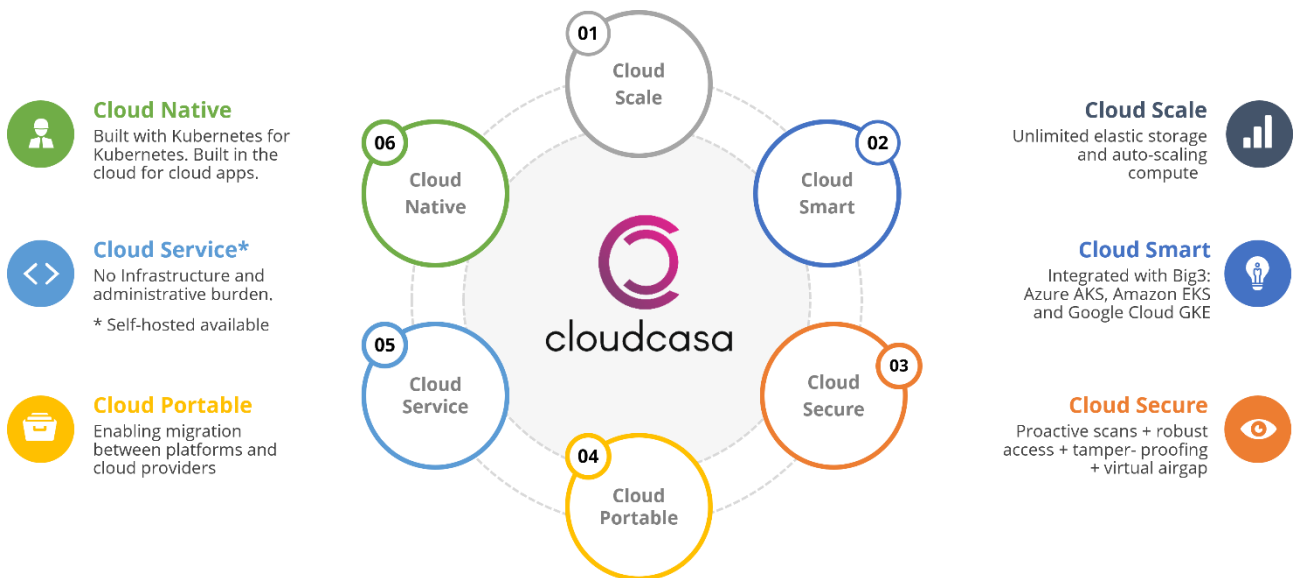
Cloud Secure: Robust protection against fraudulent system access; enables tamper-proof backups and provides a virtual air gap for backup storage.

Cloud Mobile: Migrates data across multi-cloud and multi-cluster environments with mapping of storage classes, namespaces, IAMs, and security groups.

Cloud Smart: Built in customizable Application Templates for popular databases such as MySQL, Mongo, and PostgreSQL.

Why CloudCasa vs Alternatives?

Kubernetes developers and platform teams have several options to choose from. Velero is a powerful open-source framework to protect Kubernetes data one cluster at a time, which relies on the ecosystem partners to fill gaps for multi-cluster, multi-cloud environments. Kasten by Veeam offers many of the same benefits, however it is architected as a traditional standalone application that is deployed and maintained by the users themselves.



The following are some key considerations to choose CloudCasa over alternatives.

“Full Stack” Cluster Restore for Disaster Recovery

As user adoption of Kubernetes continues to rise, more stateful workloads are running on Kubernetes. If a local outage occurs and you lose a cluster or if an entire site is destroyed, the application and data need to be quickly recovered for business continuity. For users to successfully recover missing or corrupted data, there must be a destination for the cluster to restore to. That is why most customers utilize stand-by clusters to act as the destination for recovery in these scenarios. The problem with stand-by clusters is that they need to be setup before the need for recovery and they consume cloud resources while they sit idle waiting to be used.

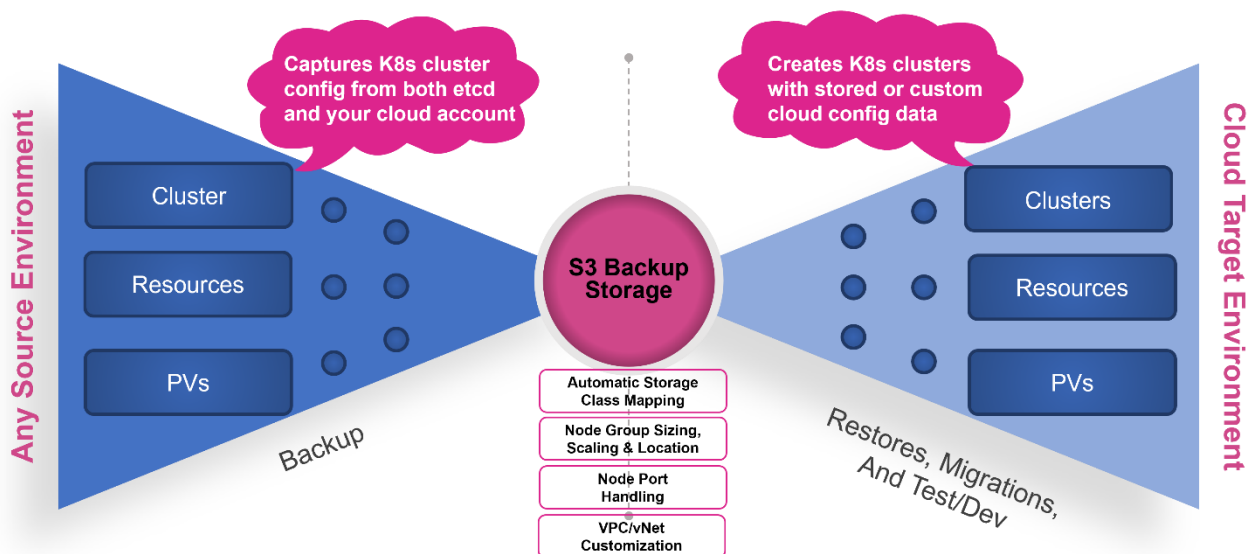
With cloud integration of CloudCasa, a backup takes a snapshot of the cloud account configuration in addition to the Kubernetes resources and data. This provides CloudCasa with the ability to spin up a brand new AKS, EKS or GKE cluster, using the cloud configuration information, in real-time during a restore. This is called Bare-Metal Cluster Recovery, and it eliminates the need for stand-by clusters.

Cross-Cluster and Cross-Cloud Data Migration

Data restores are not only used for disaster recovery but many other use cases as well. Data migration via recovery of data can also be used in the context of application transformation, application migration, and even separating staging from production in different cloud accounts. Each of these scenarios require data to move from one Kubernetes cluster, or one cloud account, to another. CloudCasa’s cross-cluster and cross-cloud restores allow users to more effectively utilize the data that they have protected.

CloudCasa Any2Cloud Recovery Bridges the Cloud Gap

Capture Current Cluster Config Data and Map it to your Target Cloud Environment



Backup-as-a-Service vs Self-Managed Backups

Updating and patching a backup application are tasks that the teams do not have to be doing. Sizing, scaling, monitoring catalog space usage etc. are additional complex tasks that are

tapping into the same resource-constrained teams running your critical business applications. CloudCasa takes all of this burden away from your teams by delivering Backup as a Service.

Virtual Air Gap vs Same Failure Domain

A standalone backup application is typically installed on one of your Kubernetes clusters being protected and if that cluster were to experience a failure, your entire backup infrastructure is down. CloudCasa runs on a highly available and well-maintained public cloud infrastructure that guarantees a 99.95% SLA. Since the chances of the service provider and customer environments are attacked at the same time is minimal, it presents customers with a virtual air gap domain.

Proactive Protection vs Last Line of Defense

All Kubernetes solutions support backup and recovery of Kubernetes workloads to S3 Storage. Most support immutability or tamper-proof capability that acts as a last line of defense. CloudCasa also has the ability to do proactive vulnerability assessments on your Kubernetes Clusters and alert you to address problems before they are exploited.

Utility Pricing vs Unpredictable Pricing

Most Kubernetes backup products charge based on worker nodes because storage allocated per worker node is still low. Worker nodes are also the primary variable in auto-scaling configurations, thereby introducing unpredictability in pricing. CloudCasa supports your long term growth of Kubernetes by employing utility-based pricing based on capacity protected. You can use unlimited worker nodes, clusters, distributions and cloud providers as long as you are under the capacity licensed.

Cloud Integrated vs Cloud Compatible

Most Kubernetes backup solutions can protect AKS and EKS Clusters and use Azure or Amazon S3 as backup destinations. CloudCasa is deeply integrated with these cloud providers - unlike alternatives - so an Azure or Amazon user can scan in their cloud accounts and auto-discover all clusters for centralized data management.

CloudCasa can perform continuous vulnerability assessments on your Azure and Amazon accounts covering over 20 popular cloud services similar to the assessments done on Kubernetes clusters. CloudCasa collects AKS and EKS metadata during cluster backup and auto-creates a new cluster based on these settings during recovery, removing the need for expensive stand-by clusters for recovery.

Your Next Step

CloudCasa is the only data protection service to provide fully automated cross-platform, in-cloud, cluster recovery irrespective of where your application containers and clusters are hosted. You can start using CloudCasa with the [Free Service plan](#), which allows you to back up your cluster resource data and protect your persistent volumes with local snapshots. No payment information is required, no limits are placed on the number of managed snapshots, worker nodes, or clusters, and retention times of up to 30 days are allowed. CloudCasa is backed by [Catalogic Software](#), a proven data protection company with over 20 years of experience safeguarding enterprise data.