DATA SHEET

cloudcasa

CATALOGIC

CloudCasa™ – Kubernetes and Cloud Database Protection as a Service

A cloud native data protection service for cloud native applications, brought to you by Catalogic, the smart data protection company.

www.cloudcasa.io

Introducing CloudCasa

CloudCasa is an easy-to-use backup service built for protecting Kubernetes, cloud databases, and cloud native applications. As a SaaS solution, CloudCasa removes the complexity of managing traditional backup infrastructure, while providing the same level of application-consistent data protection and disaster recovery that more traditional backup solutions provide for server-based applications.

With CloudCasa, your IT department doesn't need to be Kubernetes experts and your DevOps team doesn't need to be storage or data protection experts in order to protect your Kubernetes clusters and applications. CloudCasa was built as a cloud native service to support best practices for data protection and recovery for cloud native applications, and to bridge the data management and protection gap between DevOps and IT Operations.

CloudCasa Highlights

- No hardware or infrastructure to install and maintain
- No hassle, and no backup expertise needed
- Back up on-prem and in cloud
- Protects against logical, physical, accidental and malicious losses
- Always encrypted – during transit and at rest
- Supports all popular Kubernetes distributions including Red Hat OpenShift, SUSE Rancher, and VMware Tanzu
- Supports all popular Kubernetes cloud services including EKS, AKS, GKE and DigitalOcean
- Protects cluster resources and persistent volumes
- Protects Amazon RDS databases
- Control service decoupled from backup storage
- So easy developers won't mind doing backups!

CloudCasa Free Service Tier Includes

- Easy to use UI to quickly set up and manage backups
- Metadata backups sent to secure and durable cloud-based storage
- Unlimited CSI snapshots of Persistent Volumes (PVs)
- Multi-cluster data protection and management
- Unlimited number of clusters and worker nodes per account
- Flexible scheduling and data retention policies
- Protection of Amazon RDS databases, including scheduling and management of snapshots and snapshot copies.
- Point-in-time recovery of Amazon RDS databases.
- Data retention period of up to 30 days
- Community support with active participation from the CloudCasa team
- Upgradable to Premium service for additional capabilities

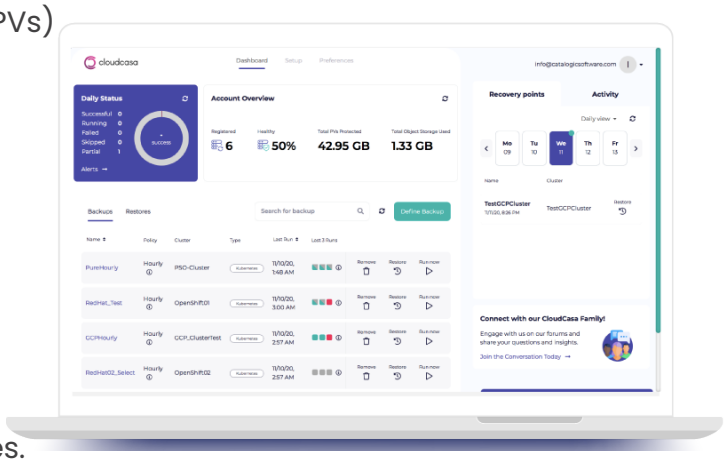

Requirements and Compatibility

Software requirements

- Kubernetes version 1.13 and higher
- PV snapshots require Kubernetes version 1.17 and higher
- Storage must use a CSI driver that supports volume snapshots at the v1beta1 API level.
  - For a list of vendors that support CSI snapshots, please see : https://kubernetes-csi.github.io/docs/drivers.html
- Supported Kubernetes distributions: Red Hat OpenShift, SUSE Rancher, and VMware Tanzu
- Supported Kubernetes cloud services: Amazon EKS, Microsoft AKS, Google GKE, IBM Cloud Kubernetes Service, and DigitalOcean.
- All Amazon RDS databases are supported, including Aurora

Note: Just because a Kubernetes distribution, cloud service, or storage device isn't listed here does not mean that CloudCasa will not work with it! We expect that nearly any variant of Kubernetes based on version 1.13 or higher will be compatible with CloudCasa.

Permissions and network requirements

The user configuring CloudCasa needs admin access to their cluster and access to the kubectl CLI. While registering your cluster in the user interface (UI), each cluster will be given a unique YAML file to be applied using kubectl.

Network access from your cluster to the CloudCasa service (agent.cloudcasa.io) on TCP port 443 is required. No ports need to be opened for inbound connections.

The login used to configure your AWS account for CloudCasa RDS backups requires administrative access, but the cross-account role created by our CloudFormation stack is limited to only the permissions necessary to perform RDS backups and restores. See our FAQ for more details.

What's Next for CloudCasa?

The CloudCasa team adopts a cloud first approach to our development process and priorities. We will soon be introducing additional features such as:

- Persistent Volume backups to secure cloud storage
- More cloud database support – Azure SQL, CloudSQL etc.
- Bring your own keys for encryption – Security is our top priority.
- Choice of backup storage regions and providers
- Replication across availability zones and increased redundancy

50 Tice Boulevard, Suite 110
Woodcliff Lake, NJ 07677 U.S.A.
201.249.8980
catalogicsoftware.com