

# Catalogic CloudCasa

## Overview

CloudCasa is Catalogic's backup and restore offering for Kubernetes environments and cloud-hosted databases. It is delivered as software-as-a-service (SaaS), and it complements Catalogic DPX, the vendor's enterprise data protection product.

CloudCasa was built on Kubernetes, and it protects a variety of Kubernetes distributions on-premises and in the cloud. CloudCasa protects the persistent container data stores as well as associated Kubernetes resources. It also protects cloud databases, including application consistent protection for Amazon RDS. Data can be replicated across clouds, migrated across accounts, and recovered to an alternate cloud zone, region, or provider.

Cloud-based storage targets are supported, including more than 20 regions in AWS and Azure as well as Backblaze B2, to enable customers to create a logical air gap in protection against ransomware. On-premises and bring-your-own object storage repositories are also supported.

CloudCasa is available through Catalogic, as well as through partners' marketplaces. To ease customer adoption, Catalogic offers a free service tier that allows customers to manage an unlimited number of snapshots with no limits on worker nodes and clusters, and backup Kubernetes resource data with a maximum retention of 30 days for both. The paid, monthly service tiers are priced based on the size of backup datasets, starting with an initial 1 TB capacity plan with unlimited retention and no limits on nodes or clusters.

## Highlights

- Backup and recovery for Kubernetes (CSI, Amazon EBS, Microsoft AKS) and cloud databases.
- Flexibility in data replication, migration, and recovery.
- Centralized cataloging, oversight and management of snapshots and backups.
- Logical air-gapped backup storage with recovery points that can be designated as immutable
- Pricing plans based on size of backup data with a free service tier for snapshots and resource data.

## Usage and Deployment

Catalogic CloudCasa is a SaaS offering specifically created for backing up and recovering Kubernetes containerized environments, and cloud-based databases.

- Characteristics
  - Backs up Kubernetes metadata and resource data and persistent volumes to cloud-based data stores.
    - Immutable “SafeLock” recover points
      - Set by GUI or an API call
  - Cross-cluster, -region, -account, -cloud provider data migration and replication
    - Based on storage class remapping
    - Facilitates disaster recovery, cluster migration, dev/test use cases
    - Minimizes egress fees within an account and supports cyber resiliency
  - Snapshot orchestration that is free forever
  - Pre- and post-backup and restore application hooks
  - APIs available with API keys
  - Encryption (data in flight and at rest) and multi-factor authentication
  - Security scanning for Kubernetes clusters and AWS cloud accounts
    - Looks for cluster, cloud, container, and network misconfiguration
- Applications/Use Cases
  - Takes snapshots and backups, allows for operational recovery and data reuse for Dev/Test
    - Multi-cluster, across hybrid and multi-cloud environments
    - Recover to alternate zone, region, provider
- Protected Environments
  - Kubernetes environments
    - Persistent data
      - CSI snapshots
      - Amazon Elastic Block Storage (EBS) persistent volume snapshots for the Amazon Elastic Kubernetes Service (Amazon EKS)
      - Microsoft Azure Kubernetes Service (AKS)
      - Block-level incremental backup of persistent volumes with only changed blocks moved
      - Rancher Longhorn
    - Cluster resources
  - Cloud databases
    - AWS RDS
- Supported Storage Targets
  - Two options:

- Catalogic provisions and manages AWS, Azure or Backblaze storage target
      - The customer does not need to be a customer of AWS the cloud provider and is not separately billed
    - Customer brings their own on-premises, public cloud or on-premises object storage target
  - More than 20 global AWS and Azure regions
  - Backblaze B2
  - DigitalOcean Spaces
  - Google Cloud Storage
  - Microsoft Azure Blob
  - Wasabi
- Deployment and Administration
  - SaaS delivery
    - Available via Catalogic, and the DigitalOcean and SUSE Rancher marketplaces
      - AWS, Azure and Red Hat marketplaces are on the roadmap.
  - Web-based GUI and APIs
  - Automated, policy-driven backup jobs and retention policies
  - Backup job monitoring and failure alerting
  - Auto-discover and import for Amazon EKS Clusters and Amazon RDS Snapshots
  - Role-based access control (RBAC)
  - Multi-tenancy and self-service
    - Clastix Capsule

## Evaluator Group EvaluScale™: Catalogic CloudCasa

Evaluator Group product review methodology “EvaluScale” assesses each product within a specific technology area. The definitions of the criteria and explanations of how products are reviewed can be found in the [Evaluation Guides](#).

	Criteria	Description	Requirement	EG View of Catalogic CloudCasa	Explanation for Catalogic CloudCasa
1	<b>Data Sources</b>	Variety and comprehensiveness of data sources that can be protected (on- and off-prem. business apps and databases, cloud-hosted IaaS, physical systems, VMs on-	<b>Meets requirements:</b> Supports two of the sources identified in “description” with the ability to automatically apply protection policies consistently. <b>Exceeds requirements:</b> Supports additional sources identified in “description” with the	<b>Exceeds requirements</b>	Supports AWS RDS as well as Kubernetes environments. Can apply protection policies consistently across environments.

		prem. and in the cloud, containers).	ability to automatically apply protection policies consistently.		
2	<b>Protection Targets</b>	Variety and comprehensiveness of locations for storing backup and retention data, to meet a variety of service level, compliance, and cyber security requirements.	<b>Meets requirements:</b> Provides operational recovery and uses some form of isolation/air gapping as well as immutability to satisfy “3-2-1” best practices. <b>Exceeds requirements:</b> Supports on-prem. and cloud-based storage targets, replication, and tiering to archive storage.	<b>Exceeds requirements</b>	Supports Catalogic-provided cloud-based storage targets, as well as bring-your-own on-prem. and cloud object storage targets. Facilitates logical air gapping, offers control of immutability of underlying storage via GUI or API. Data can be replicated across clouds, migrated across accounts, and recovered to an alternate cloud zone, region, or provider.
3	<b>Backup Control and Oversight</b>	Policies governing data backup jobs.	<b>Meets requirements:</b> Control over backup frequency and target. Allows for the same protection policies to be applied irrespective of the source location. <b>Exceeds requirements:</b> Can identify and triage failures. Protection policies can account for security, data governance and compliance requirements.	<b>Meets requirements</b>	Supports snapshots, and backups with customer-controlled frequency as well as target (either Catalogic- or customer-provided). Policies can be applied globally. Monitors backup jobs and identifies failures.
4	<b>Recovery Flexibility</b>	Comprehensiveness of recovery options.	<b>Meets requirements:</b> Operational recovery to the cloud. <b>Exceeds requirements:</b> Self-service portal for restores. Cross-region replication. Local cache for faster recovery. Granular recovery options (e.g., can recover an entire VM or individual files).	<b>Meets requirements</b>	Data can be replicated across clouds, migrated across accounts, and recovered to an alternate cloud zone, region, or provider (including to an on-prem. Kubernetes environment). Supports full and partial restores. Self-service portal. No local cache.
5	<b>Data Management</b>	Data lifecycle management (creation through deletion, including data migration).	<b>Meets requirements:</b> Catalogs all data discovered to be able to manage the protected copies and identify unprotected data. Can search for data and identify different file versions. Monitoring and logging of, as well as reporting on, backup activities.	<b>Exceeds requirements</b>	Centralized cataloging, oversight and management of snapshots and backups. Searching, monitoring, and reporting – as well as data migration and automated tiering for long-term retention.

			<b>Exceeds requirements:</b> Data migration. Automated tiering to archive storage. APIs/application hooks.		
6	<b>Economics</b>	Spans licensing details, scaling of resources, as well as the ability to help customers calculate TCO.	<b>Judgement-based, with evaluation criteria including:</b> TCO evaluation tool. Pricing includes storage target capacity. Capabilities to avoid egress charges. Automatic scaling based on requirements. Charging based on the data protected, not the infrastructure running. Option for a free tier (at least for trial purposes).	<b>Exceeds requirements</b>	Pricing plans based on size of backup data with a free service tier for snapshots and resource data. Pricing includes storage target capacity. Egress fees apply only with cross-account replication. Auto-scaling of resources.
7	<b>Regulatory Compliance</b>	Meets regulatory requirements for specific types of data.	<b>Varies based on individual needs; basics include: Meets requirements:</b> Provides immutability and the ability to isolate a data copy (e.g. air gapping). RBAC, encryption, encryption key management. <b>Exceeds requirements:</b> Provides identification, and centralized oversight and analytics (e.g. identifying PII). Provides dashboarding, reporting and AI/automation to ensure and prove compliance. Offers a dedicated environment. AI for identification of security vulnerabilities.	<b>Meets requirements</b>	Immutability (“SafeLock,” controlled via GUI or an API call). Isolated data copies in the cloud. RBAC. Encryption of data at rest and in flight.
8	<b>Security</b>	Provides defense against cyberthreats, including (and especially) ransomware.	<b>Meets requirements:</b> Air gapping and immutability. RBAC. Encryption and encryption key management. <b>Exceeds requirements:</b> Recovery into a sandbox environment. Automation and other features for accelerated recovery. AI for threat detection.	<b>Meets requirements</b>	Immutability (“SafeLock,” controlled via GUI or an API call). Isolated data copies in the cloud. RBAC. Encryption of data at rest and in flight.
9	<b>Management and Manageability</b>	Simplicity of day-to-day management. Migrations, integrations, and training required.	<b>Meets requirements:</b> Automated discovery of resources. Automated backup jobs. Backup job monitoring and alerting of failures.	<b>Meets requirements</b>	Automated resource discovery. Automated backup jobs. Monitoring and alerting of failures. Option for Catalogic to provision and manage the

			<b>Exceeds requirements:</b> The service provider handles sizing and upgrading of software components. 24/7 management, monitoring of environment for adherence with SLAs. Remedial actions/compensation if SLAs are not met. API support. Allows data to be transitioned if the service is discontinued.		storage target; the customer does not need to be a customer of AWS or Azure and in this model is not separately billed. API support.
<b>10</b>	<b>Vendor Support / Product Viability</b>	Long-term support and advancement of the product from the vendor.	<b>Judgement-based.</b> "Meets Requirements" typically requires more than five years of company existence and product general availability (GA); "Exceeds" typically requires more than 15 years.	<b>Meets requirements</b>	Catalogic was spun out of Syncsort in 2013; CloudCasa was launched in November 2020.

## Evaluator Group Opinion: Differentiating Elements for Catalogic CloudCasa

While other container data protection solutions are self-hosted, have been retrofitted from a pre-existing software, or are bundled with the container storage itself, CloudCasa is simultaneously stand-alone, SaaS-delivered and built on Kubernetes. Additionally, CloudCasa benefits from Catalogic’s background in snapshot and copy data management, offering customers a single tool for centralized cataloging and monitoring of, as well as reporting on, snapshots and backups. It also eliminates the need for separate tools for EBS and CSI snapshots and backups, and RDS protection. CloudCasa’s ability to auto-create clusters during recovery based on the configuration of the backed-up cluster shortens time-to-recovery. Its ability to review the environment’s security posture can help to thwart ransomware attacks by uncovering vulnerabilities. The free service tier is a good opportunity for customers to trial CloudCasa. The capacity-based pricing model is intuitive and fair in that customers are charged for data they are protecting, versus per node and cluster models that charge for what they are running.

The sale of its ECX copy data management product to IBM in May 2021 will allow Catalogic to sharpen and accelerate its product development and go-to-market initiatives pertaining to CloudCasa – for instance, offering pre-built application hook templates for additional databases, availability in additional catalogs and marketplaces, and further enhancing its security capabilities and services.

Information that is more detailed is available at <http://evaluatorgroup.com>

Copyright 2022 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.

Eval(u)Scale™ is a methodology developed by Evaluator Group to determine the value offered by a product within a technology area. Evaluation criteria are developed through IT client engagements and an understanding system usage. The definitions of the criteria and explanations of how products are reviewed can be found in the Evaluator Group Evaluation Guide, which is reflected in the Eval(u)Scale published in individual Product Briefs. The Eval(u)Scale for each product is determined through in-depth review and analysis attained from vendor interviews, reviews of user / administration guides, hands-on testing and/or client engagements.