

# Robin Cloud Native Storage

## Container-Native Storage

January 2022

---



**Evaluator Group**

*Making the best decisions for Information Management*

# Robin Cloud Native Storage

## Overview

Robin Cloud Native Storage provides block and file storage that runs on any Kubernetes platform.

Startup Robin.io first launched CNS in 2019. Although Cloud Native Storage is sold as an independent product, it is also bundled with the Robin Cloud Native Platform, which also includes virtual networking, and application workflow automation built on Kubernetes. Robin.io also sells a multi-data center orchestration and automation platform (MDCAP).

Robin.io supports Red Hat OpenShift, Rancher Kubernetes Engine, managed Kubernetes such as Azure AKS, AWS EKS, Google GKE, and IBM IKS, and on-prem Google Anthos. It installs natively as a Kubernetes operator.

Container Native Storage (CNS) is a category of software-defined storage (SDS) that's embedded in a container cluster, providing a virtualized pool of storage using devices internal to each host. This is a hyperconverged infrastructure (HCI) cluster except that the software runs in a container, instead of as a VM. CNS in a container orchestration platform (like Kubernetes) is somewhat analogous to VMware vSAN running as a service in vSphere, creating a scale-out compute environment.

The benefits of Container-Native Storage with Kubernetes are also similar to those of HCI. A CNS eliminates the need for external storage systems to support containers and can scale with the Kubernetes cluster. As a storage layer, CNS is integrated into Kubernetes, simplifying management compared with an external storage system.

Robin Cloud Native Storage runs on-premises on bare metal or virtual machines, or in public clouds. Robin Cloud Native Storage features include disk discovery and pooling, synchronous replication, encryption (but no built-in key management), compression, thin provisioning, capacity management and forecasting, IO hotspot reduction and protection against disk and node failures.

Robin.io's data management features include:

- application-consistent snapshots and clones
- copy data management

### Highlights

- Storage features include disk pooling, replication, encryption and thin provisioning
- Performs app-consistent snapshots, backup and recovery
- Can run as part of Robin Container Native Platform built on Kubernetes

- application-based backup and recovery
- backup and snapshot scheduling with retention policies and quality of service.

Application-consistent snapshots and clones capture entire applications, including storage volumes, configurations and metadata.

For disaster recovery, Robin Cloud Native Storage can move entire Kubernetes apps with their data among on-prem, public and hybrid cloud environments across clusters to other on-premises and cloud locations. Robin Cloud Native Storage provides replication at the node, rack, data center and zone levels. Robin.io's Container Storage Interface (CSI) exposes storage to workloads running on Kubernetes. Robin.io's CSI driver includes three storage classes used to provision CNS storage:

- Default StorageClass has no features enabled, and can be used for standard ReadWriteOnce (RWO) volumes for block storage and ReadWriteMany (RWX) volumes for files.
- "Repli-3" StorageClass includes three replicas
- "Immediate" StorageClass creates a volume as soon as a Persistent Volume Claim (PVC) is created rather than waiting for the first consumer of that volume.

The Robin.io advanced CSI Storage Class provides these storage management features:

- Disk discovery and pooling
- Replication
- Encryption
- Compression
- Thin provisioning
- Capacity management/forecasting
- IO hotspot reduction
- Protection against disk/server failures

Robin Cloud Native Storage protects data on nodes and disks through volume level synchronous replication and re-sync on failures. Robin.io supports two replicas maintained on different disks to protect against a single disk or node failure, or three replicas to protect against two disk/node failures. If a replica disk becomes unavailable, Robin Cloud Native Storage evicts it from the group and writes to the healthy replicas. When the failed disk is available, Robin.IO syncs it with other healthy disk before resuming writes. If a serious disk or node error occurs, Robin Cloud Native Storage issues an alert for the storage admin, who can add a new disk if needed.

## Usage and Deployment

Robin.io positions its software as application-aware, meaning it tunes performance based on the application type. Robin.IO installs on any Kubernetes cluster, and its operator discovers the underlying media types.

It uses this information for advanced data placement. For instance, if an application requires high performance, Robin Cloud Native Storage will place the compute and storage on the same host. If an application uses shards and each shard has a replica, Cloud Native Storage will place the replicas on different physical hosts. Robin.io will identify IOPS hotspots and move volumes to different disks, although that data movement requires manual intervention.

Robin Container Native Storage is primary and secondary storage. Its built-in data protection enables backups to cloud and remote clusters, although customers can use independent backup products for data protection with Robin.io's primary storage.

Robin Express is a full-featured free version of CNS that is limited to five nodes and 5 TB of data, and community support. Robin Enterprise is the licensed version of Cloud Native Storage with unlimited node and storage capacity, enterprise support and per-node pricing. Robin Enterprise is sold through pay-as-you-go pricing on Red Hat marketplace, at \$0.42 per node per hour.

Robin.IO has Cloud Native Storage partnerships with Google and IBM. Google resells Robin storage for Google Anthos deployments on bare metal. Robin can also be used as storage with IBM Cloud Pak for Data, IBM's analytics platform for AI that runs on OpenShift.

## Evaluator Group EvaluScale™: Robin Cloud Native Storage

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	EG View of Robin Cloud Native Storage	Explanation for Robin Cloud Native Storage	
1	Scale performance/capacity	Ability to easily scale capacity, nodes, clusters and performance; storage media/cloud support	Meets requirements	5,000 nodes per cluster; more than 100,000 volumes per Robin Platform; NVMe, HDD, DRAM cache
2	Enterprise Storage Features	Provide unique or advanced efficiency, management, automation outside of Kubernetes	Exceeds Requirements	Read-on-write snapshots, clones, QoS, replication, tiering, thin provisioning, compression, auto-rebuild with changed-block tracking and volume relocation in event of pod or node failure.
3	Resiliency	Features to keep data available during node failures	Exceeds requirements	RAID 6, up to 3 replicas
4	Security	Encryption for data at rest and in transit, key management, role-based access controls	Meets requirements	Encryption, RBAC
5	Storage Management/Monitoring	Integration of storage management, analytics into Kubernetes and other container management platforms, monitoring tools	Meets requirements	Tight Kubernetes integration; advanced CSI; lacks predictive analytics
6	Ease of Deployment	Installs on Kubernetes, container platforms on-prem and in cloud	Meets requirements	Installs as container, may be packaged with Robin Platform K8s distro; lacks predictive analytics
7	Storage Access Methods	Type of storage protocols supported in containers	Meets requirements	Native support for block with file system overlay, no object storage support

8	Vendor / Product Stability	Company viability, product longevity, impact of OEM partnerships	Area for development	Small company, relies on VC funding, has affinity among DevOps community; Google and IBM partnerships
9	Cloud Options	Supports public cloud deployments for hybrid clouds	Meets requirements	Supports all major public clouds
10	Economics	Effective CAPEX and OPEX, including data reduction, licensing options	Area for development	No dedupe, storage space reclamation

## Evaluator Group Opinion: Differentiating elements for Robin Cloud Native Storage

Robin.io is among a small group of vendors that started specifically to develop storage for containers and cloud-native applications. Robin Cloud Native Storage has a set of enterprise features, including quality of service for storage, compute and memory resources, several replication modes and auto-resync. Still a relatively new product, Robin.io lacks data deduplication and advanced analytics.

Traditional storage solutions can meet these data services requirements but were designed for VM environments, not for the scale that container-based applications run. Products such as Robin Cloud Native Storage bring management to the container level.

Robin.io is a small company that has only been in the market for a few years, and depends on venture funding. Its software was created for developers and its platform can establish a following among DevOps but may need partnerships with established storage vendors as well as networking companies to fulfill its long-term strategy of enabling telcos to build 5G networks.

More information 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.

EvaluScale™ 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 EvaluScale published in individual Product Briefs. The EvaluScale 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.