

ROBIN Storage - Advanced Data Management for Kubernetes

Use the preferred storage for GKE to run mission-critical stateful applications on your Kubernetes. Bring advanced data management to Kubernetes-native frameworks including Helm and Operators.

Highlights

- » **Get high performance** enterprise-grade storage trusted and validated by Google
- » **Quickly and easily deploy** enterprise workloads on any Kubernetes distribution
- » **Protect app+data** with replication, snapshots, backup & recovery, and enterprise-grade security
- » **Bring automated management** of app+data (not just storage) to kubectl, Helm, and Operators
- » **Guarantee QoS** for high priority applications by setting IOPS limits per application
- » **Collaborate and share** across geos and teams by cloning app+data in minutes
- » **Easily move app+data**, between on-prem and cloud(s)

The Need For Data Management

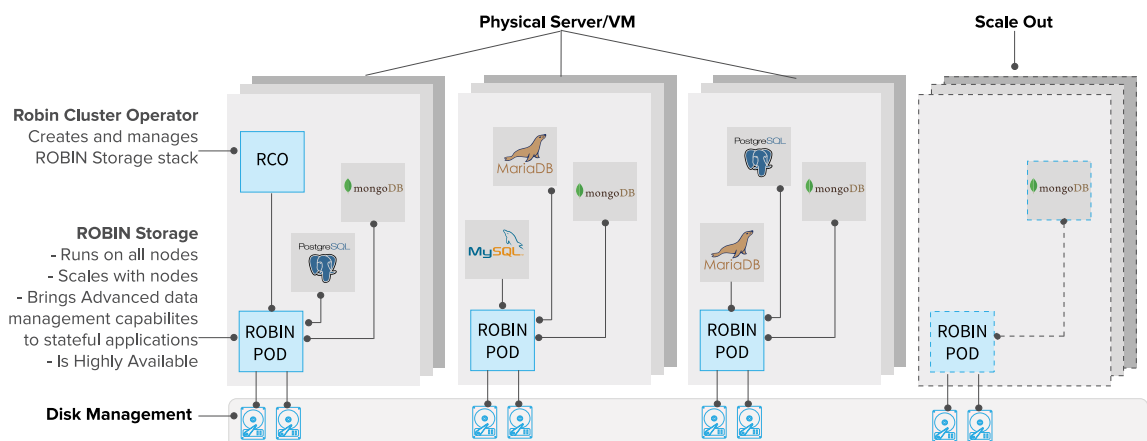
Stateful applications such as PostgreSQL, MySQL, MongoDB, Elastic Stack, Kafka, and MariaDB require advanced data management capabilities in order to:

- » **Release products on schedule:** Automated lifecycle management for app+data (not just the storage) is required to save valuable time at each stage of the lifecycle. Multiple teams (Dev/Test/Ops) need a mechanism to collaborate without procedural delays.
- » **Recover from system failures:** App+data protection capabilities such as point-in-time snapshots, backup, and restore are required to recover from system failures and user errors.
- » **Avoid infrastructure lock-in:** The ability to migrate from on-prem to cloud and vice versa, and among the public clouds is needed to avoid infrastructure lock-in.
- » **Deliver predictable performance:** To guarantee QoS and to ensure high priority applications do not miss SLAs, you need the ability to set IOPS limits per app.
- » **Eliminate security vulnerabilities:** Enterprise-grade security is required with authentication and encryption to ensure your data is safe.

ROBIN Storage - Advanced Data Management for Kubernetes

ROBIN Storage is a purpose-built container-native storage solution that brings advanced data management capabilities to Kubernetes. It provides automated provisioning, point-in-time snapshots, backup and recovery, application cloning, QoS guarantee, and multi-cloud migration for stateful applications on Kubernetes.

ROBIN Storage for Kubernetes - GKE



ROBIN Storage - Benefits and Business Impact

DELIVER PRODUCTS FASTER

Automate application management

ROBIN Storage is application-aware. All lifecycle operations are performed in the context of the app+data, not just the storage volumes. For example, when you snapshot or backup a stateful application, ROBIN Storage includes data as well as the application configuration and metadata. Automating the apps+data management saves valuable time at each lifecycle operation.

Collaborate faster across teams

With ROBIN Storage, you can clone entire application environments, and share them across teams. While developing new features, your Dev team can use a clone of the production environment as the starting point to minimize errors. Upon finishing the implementation, the Dev team can create a clone of their environment and handover to QA for testing.

Use Kubernetes-native frameworks

You can provision and manage stateful applications using Kubectl, Helm, and Operator frameworks with ROBIN Storage. Using standard frameworks helps developers be productive.

RUN MISSION-CRITICAL APPS CONFIDENTLY

Data protection for always-on applications

ROBIN provides built-in data replication to ensure High Availability. Point-in-time snapshots for apps+data enable time-travel and recovering from user mistakes, and apps+data Backup and Restore functionality enables you to recover from application crashes and/or data center downtime (Disaster Recovery).

Experience bare-metal performance

ROBIN Storage delivers near bare-metal performance with the flexibility of software defined storage. The superior performance enables you to run mission-critical applications that require high performance storage.

Never miss application SLAs

High priority applications need predictable performance. You can guarantee QoS for high priority applications by setting maximum IOPS limits per application. This enables you to consolidate applications on shared infrastructure and reduce hardware cost, without worrying about the noisy neighbors issue.

Rest assured your data is secure

ROBIN Storage authenticates user's Kubernetes-native mechanism. Your data is encrypted at rest and in transit to ensure it is protected against unauthorized intruders. ROBIN storage also comes equipped with check-sum algorithms to make your data tamper-proof.

FUTURE-PROOF YOUR ENTERPRISE

Cloud-native data management

Kubernetes is rapidly becoming the de-facto platform to manage applications. Cloud-native architecture powered by Kubernetes, containers, and ROBIN Storage enables you to modernize your application and data infrastructure and move towards the future.

Get Hybrid and Multi-cloud flexibility

You can easily move apps+data, between on-prem and multiple clouds. Enjoy cost-efficiencies by choosing the most economic infrastructure at any given point. Avoid infrastructure lock-in with the flexibility to run your applications where you want.

Avoid vendor lock-in

Kubernetes-based architecture gives you complete control of your infrastructure. With the freedom to move your workloads across private and public clouds, you avoid vendor lock-in. ROBIN Storage also provides you the flexibility to leverage your existing SAN/NAS/DAS.

To learn more and to try ROBIN visit: robin.io