# Robin Hyper-Converged Kubernetes Platform

Extend Kubernetes for data-intensive applications such as Hortonworks, Cloudera, Elastic stack, RDBMS, NoSQL, and AI/ML applications

## What is ROBIN?

ROBIN is the industry's first hyper-converged Kubernetes platform for big data, databases, and AI/ML that provides a self-service App-store experience. Hyper-converged  Kubernetes is a software-defined  application orchestration framework that combines containerized storage, networking, compute (Kubernetes), and the application management layer into a single system.

Robin.io helps enterprises achieve faster roll-out of critical IT initiatives (e.g. containerization, cloud migration, multi-cloud strategy, cost-consolidation) and LoB initiatives (e.g. AI/ML, analytics projects).

This software-only solution runs on-premises in your private data center or in public-cloud (AWS, Azure, GCP) environments and enables 1-click deployment of any big data, database or AI/ML application. ROBIN enables 1-click simplicity for lifecycle management operations such as snapshot, clone, patch, upgrade, backup, restore, scale, & QoS control of the entire application. ROBIN solves fundamental challenges of running big data & databases in Kubernetes & enables deployment of an agile & flexible Kubernetes-based infrastructure for Enterprise Applications.

## Key Benefits

Leverage Kubernetes for data-heavy applications

Get agile, flexible Kubernetes-based infrastructure

Accelerate your  critical Enterprise IT and LoB initiatives

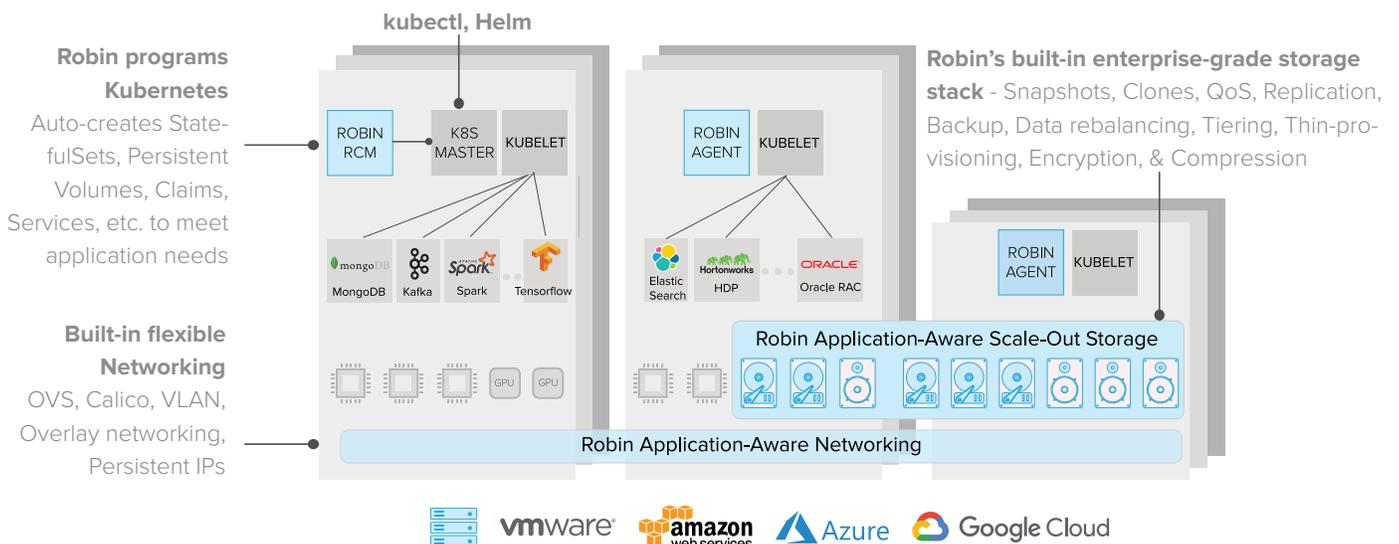## Hyper-Converged Kubernetes Stack Components

**Application Management Layer** - Manage Applications and configure Kubernetes, Storage & Networking with Application workflows.

**Kubernetes** - Run big data and databases in extended Kubernetes, eliminating limitations that restrict Kubernetes to micro-services applications.

**Built-in Storage** - Allocate storage while deploying an application or cluster, share storage among apps and users,  get SLA guarantees when consolidating, support for data locality, affinity, anti-affinity and isolation constraints, and tackle storage for applications that modify the Root filesystem.

**Built-in Networking** - Set networking options while deploying apps and clusters in Kubernetes and preserve IP addresses during restarts.

## Robin Implementation of Hyper-Converged Kubernetes

**Robin programs Kubernetes**
Auto-creates State-fulSets, Persistent Volumes, Claims, Services, etc. to meet application needs

**Built-in flexible Networking**
OVS, Calico, VLAN, Overlay networking, Persistent IPs

**Robin's built-in enterprise-grade storage stack** - Snapshots, Clones, QoS, Replication, Backup, Data rebalancing, Tiering, Thin-provisioning, Encryption, & Compression

kubectl, Helm

# Robin Hyper-Converged Kubernetes Platform Features and Benefits.

| Features | Benefits |
|---|---|
| **Rapid Deployment** - Self-service 1-click App-store experience. | Slash deployment and management times from weeks and hours to minutes. Deploy and manage data-heavy apps and services in Kubernetes. |
| **Control QoS** - Dynamic control QoS for every resource - CPU, Memory, Network and Storage. | Get complete visibility into the underlying infrastructure, set min and max IOPs, eliminate noisy neighbor issue, and gain performance guarantee. |
| **Rapid clones** - Clone the entire application along with its data - thick, thin, or deferred. | No performance penalties, backup data with ease, share data among users and applications, among dev, test, and prod, with no additional storage. |
| **Application Snapshots** - Take unlimited full application cluster snapshots, which include application configuration + data | Restore or refresh a cluster to any point-in-time using snapshots. Roll back easily with 1-click to the last snapshot in case of data corruption. |
| **Scale** - Decouple compute and storage, scale independently. | Scale out - add nodes. Scale up - increase CPU, Memory and IOPs. |
| **High Availability** - No single point of failure - get reliable crossover and detect failures. | Get automatic App-aware data failover for complex distributed applications on bare metal - ROBIN is the ONLY product to provide HA for apps that persist state inside Docker images. |
| **Upgrade -** Automated rolling upgrade of application containers that is integrated with CI/CD pipeline. | Safe-Upgrade technology guarantees that failed upgrades can be rolled back without disrupting the application. |

## Enterprise Data Apps-as-a-Service - Sample Customer Deployments

| | | |
|---|---|---|
| Fortune 500 Financial Services Leader<br><br>**11 billion** security events ingested and analyzed in a day<br><br>DevOps simplicity for Elasticsearch, Logstash, Kibana, Kafka | Global Networking and Security Leader<br><br>**6 Petabytes** under active management in a single Robin cluster<br><br>Agility, consolidation for Cloudera, Impala, Kafka, Druid | Global Technology Company - Travel Industry<br><br>**400 Oracle RAC** databases managed by a single Robin cluster<br><br>Self-service environment for Oracle, Oracle RAC |

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