

Elastic Stack-as-a-Service with Kubernetes

Turbocharge your Elastic Stack infrastructure using cloud-native architecture. Improve the agility and efficiency of your Developers, Operation teams, and Data Scientists.

Highlights

- » **Provision** custom Elastic Stack clusters on Kubernetes in minutes
- » **Provide** self-service experience to improve developer and data scientist productivity
- » **Scale-up/scale-out** Data Nodes dynamically in seconds, without interrupting cluster operations
- » **Ensure** data locality for Data Nodes for better performance
- » **Secure** your data with encryption at rest and in-motion, authentication, and RBAC
- » **Ensure** high availability using rack-aware placement rules for Master and Data Nodes
- » **Consolidate** multiple ELK clusters on shared infrastructure to reduce hardware footprint
- » **Trade** resources among ELK clusters to manage surges and periodic compute requirements

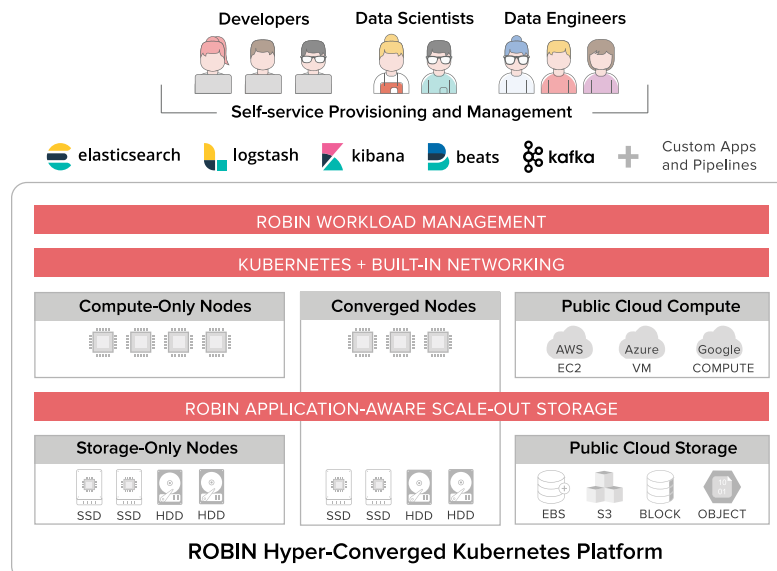
Top 5 Challenges for Elastic Stack Management

The Elastic Stack has made it simple to ingest, search, analyze, and visualize data. However, following challenges keep organizations from unlocking the full potential of the Elastic Stack:

- » **Monolithic ELK clusters:** Giant shared ELK clusters slow down developers and data scientists because they cannot make the necessary customizations (e.g. a specific versions of Elasticsearch, or Beats).
- » **Provisioning custom ELK clusters:** Developers need custom stacks with different versions and combinations of Elasticsearch, Logstash, Kibana, Beats, and Kafka. They cannot quickly provision/decommission these custom stacks themselves, without creating IT tickets.
- » **Dynamic scaling to meet sudden demands:** If a Data Node runs out of resources, there is no easy way to scale-up the node “on-the-fly” by adding more memory or CPU. Scaling-out to add more Data Nodes can also take weeks due to process delays.
- » **Security concerns:** Without encryption at rest, your data is vulnerable. All data should be treated as sensitive, and be secured.
- » **Multi-cluster strategy leads to massive hardware costs:** Creating dedicated clusters for individual “tenants” (teams, workloads, applications etc.) is a good strategy. But requires provisioning each cluster for peak capacity, leading to significant hardware underutilization.

ROBIN Hyper-converged Kubernetes Platform

ROBIN Hyper-converged Kubernetes platform extends Kubernetes with built-in storage, networking, and application management to deliver a production-ready solution for the Elastic Stack. ROBIN automates the provisioning and management of big data clusters to deliver an “as-a-service” experience with 1-click simplicity to data engineers, data scientists, and developers.



Solution Benefits and Business Impact

ROBIN brings together the simplicity of hyper-convergence and the agility of Kubernetes.

DELIVER INSIGHTS FASTER

Self-service experience

ROBIN provides self-service provisioning and management capabilities to developers, operations teams, and data scientists, significantly improving their productivity. It saves valuable time at each stage of application lifecycle.

Provision custom stacks in minutes

ROBIN has automated the end-to-end cluster provisioning process for the Elastic Stack, including custom stacks with different versions and combinations of Elasticsearch, Logstash, Kibana, Beats, and Kafka. The entire provisioning process takes only a few minutes.

Eliminate “right-size” planning delays

DevOps and IT teams can start with small deployments, and as applications grow, they can add more resources. ROBIN runs on commodity hardware, making it easy to scale-out by adding commodity servers to existing deployments.

Scale on-demand

No need to create IT tickets wait for days to scale-up Data Nodes by adding more memory, CPU, or Storage, or to scale-out by adding more Data Nodes. Cut the response time to few minutes with 1-click scale-up and scale-out.

REDUCE COSTS

Improve hardware utilization

ROBIN provides performance isolation and role-based access controls (RBAC) to consolidate multiple ELK workloads without compromising SLAs and QoS, increasing hardware utilization and reducing cost. Ensuring data locality for Data Nodes also provides better performance, reducing hardware footprint.

Simplify lifecycle operations

Native integration between Kubernetes, storage, network, and application management layer enables 1-click operations to provision, scale, snapshot, clone, backup, migrate ELK clusters, reducing the administrative cost.

Trade resources among ELK clusters

Reduce your hardware cost by sharing the compute resources between clusters. If an ELK cluster runs the majority of its batch jobs during the night-time, it can borrow a resource from an adjacent ELK cluster with day-time peaks, and vice versa.

FUTURE-PROOF YOUR ENTERPRISE

Standardize on Kubernetes

Modernize your data infrastructure using cloud-native technologies such as Kubernetes and Docker. ROBIN solves the storage and network persistency challenges in Kubernetes to enable provisioning and management of mission-critical ELK deployments. ROBIN also eliminates the need to manually define and manage individual K8S configuration objects such PVC, PV, StatefulSets, Services, etc.

No cloud lock-in

Kubernetes-based architecture gives you complete control of your infrastructure. With the multi-cloud portability, you have the freedom to move your workloads across private and public clouds, you avoid vendor lock-in.

Enterprise-grade security and HA

ROBIN provides encryption at rest out of the box, bringing an extra layer of security. Rack-aware placement rules for Master and Data Nodes ensures your HA setup is production-ready.

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