

Elastic Stack as a Service on Kubernetes

Optimize Elastic Stack using cloud-native architecture, improving the agility and efficiency of your developers, operation teams, and data scientists.

Top Four Challenges for Elastic Stack Management

Elastic Stack (Elasticsearch, Kibana, Logstash, Beats, and Kafka, also known as the ELK Stack) has made it simple to ingest, search, analyze, and visualize data. However, the following challenges keep organizations from unlocking the full potential of their Elastic Stack deployments:

Provisioning custom ELK clusters: Developers need custom stacks with different versions and combinations of ELK. 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.) can be a good strategy. But it requires provisioning each cluster for peak capacity, leading to significant hardware underutilization.

Highlights

- Provision custom Elastic Stack clusters on Kubernetes in minutes
- Provide a 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 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

Solution Benefits and Business Impact

Robin brings together the simplicity of automation 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. The platform saves valuable time at each stage of the application lifecycle.

Provision custom stacks in minutes

Robin automates 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. With Robin, provisioning is a one-click process that takes only a few minutes.

Eliminate “right-size” planning delays

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

Scale on demand

With Robin, there is 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. One-click scale up and scale out will enable you to cut response time to a few minutes.

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 the application management layer enables one-click operations to provision, scale, snapshot, clone, backup, and migrate ELK clusters, reducing the administrative cost.

Share resources among ELK clusters

Robin enables you to reduce your hardware cost by sharing compute resources between clusters. If an ELK cluster runs the majority of its batch jobs during the night, it can borrow a resource from an adjacent ELK cluster with daytime peaks, and vice versa.

Future-Proof Your Enterprise

Standardize on Kubernetes

Robin automates data infrastructure using cloud-native technologies, including Kubernetes and containers. 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 Kubernetes configuration objects, such as PVC, PV, StatefulSets, services, etc.

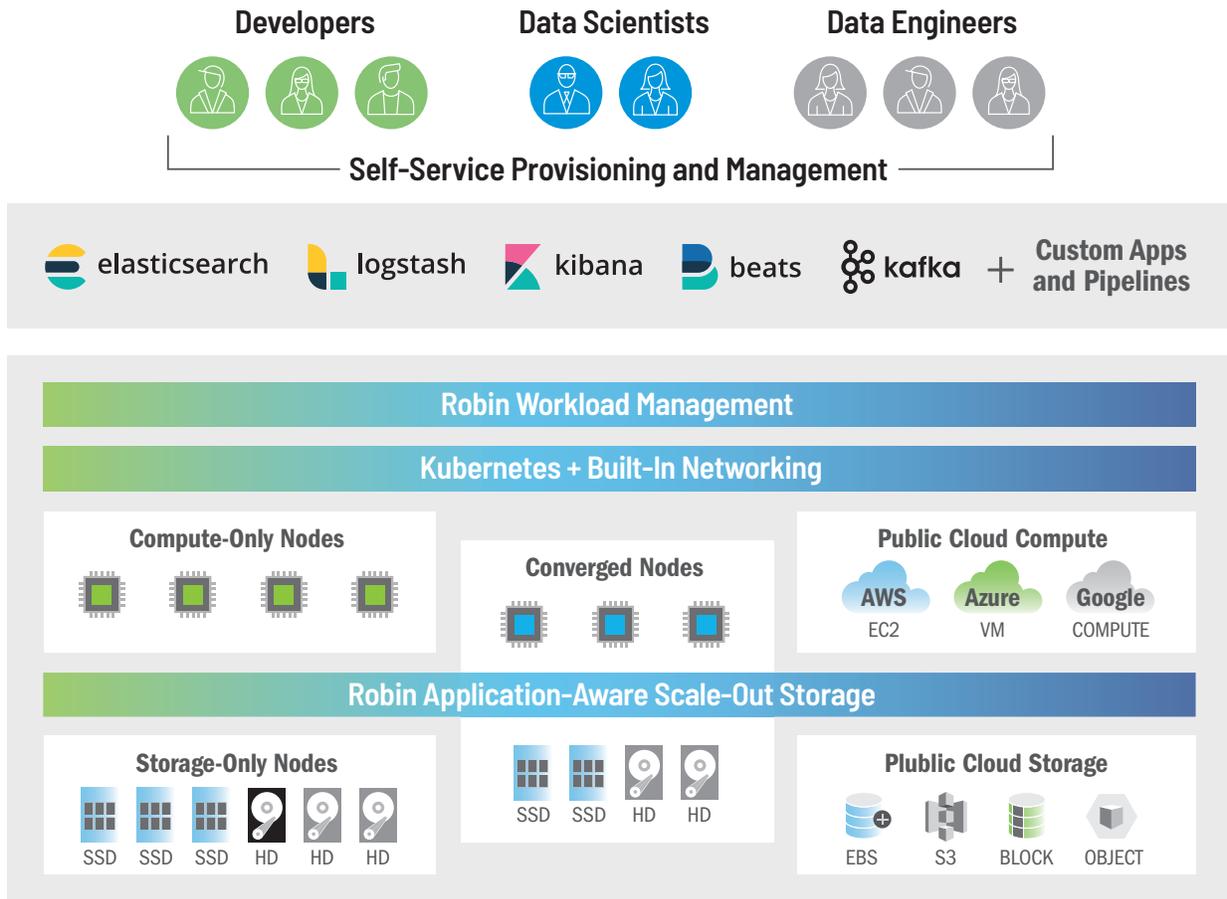
No cloud lock-in

The Robin Kubernetes-based architecture gives you complete control of your infrastructure. With its multi-cloud portability, the platform gives you the freedom to move your applications across private and public clouds and 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.

The Robin Platform Enables an “as a Service” Experience



The Robin application automation platform enables enterprises to deploy complex data- and network-centric application pipelines on Kubernetes. Robin automates provisioning and day-two operations so you can deliver a self-service experience with the simplicity of one-click deployment for developers, DBAs, and data scientists.