JioCloud

# Managed Kubernetes Services

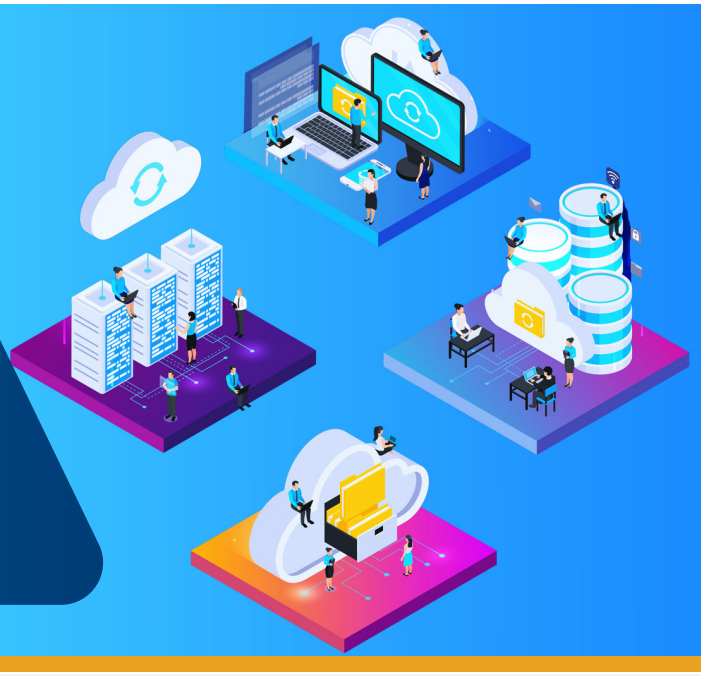

**Overview**

JioCloud Managed Kubernetes offers a fully managed enterprise-grade Kubernetes platform with the flexibility and control for modern workloads. Deploy secure and high-availability clusters with integrated CI/CD, observability and enterprise-grade policies in minutes. Choose CPU or GPU worker nodes, auto-scale based on traffic, and automatically recover from failures. Our managed service lets your teams innovate faster, while we manage the complexity.

## Key Features

- **One-click cluster provisioning**
  You can quickly deploy production clusters while we handle the underlying networking, security, and autoscaling configuration for you.
- **Flexible worker nodes**
  Optimise performance and cost by provisioning the right compute for the job. Deploy standard CPUor high-performance GPU worker nodes to match your specific application needs.
- **Elastic auto-scaling**
  Automatically scale pods and clusters based on demand spikes.
- **Automated upgrades and patching**
  You can remain compliant with zero downtime patching, and upgrades.
- **Built-in integrated observability**
  You can view observable metrics, dashboards and alerts with real-time control in HaKube.
- **Self-healing**
  Automatically replace failed nodes and restart failed pods.
- **Role-based access control (RBAC)**
  Control granular permissioning across teams.
- **Service mesh ready**
  Add secure service-to-service communication at any time with Istio integration.

JioCloud

# Benefits

- **Decreased time-to-market**
  Get production clusters up and running in minutes, not days.
- **Reduce K8s management effort by up to 60%**
  Let JioCloud manage the infrastructure and scale accordingly—our experts will work for you.
- **Improve workload performance**
  Deploy ML, HPC, or real-time apps on optimized CPU/GPU resources.
- **Stay secure and compliant**
  The solution is built with RBAC, encryption, logging, and audit-ready controls.
- **Run anywhere**
  Easily and consistently extend Kubernetes across hybrid and multi-cloud environments.

## Technologies Supported

- **Kubernetes:** CNCF conformant 1.2x series
- **Node OS:** Ubuntu / RHEL (Linux); GPU pools available (NVIDIA)
- **Container Runtime:** containerd; OCI images
- **CNI:** Calico (policy), others on request
- **Ingress:** NGINX/Envoy class controller; integration with **Jio Cloud API Gateway (APISIX)**
- **Load Balancer:** Radware (edge) and **MetalLB** (internal) options
- **Service Mesh (optional):** Istio for mTLS, traffic shaping, and telemetry
- **Storage (CSI):** OpenStack Cinder, NetApp Trident; RWX/RWO classes; snapshot/clone
- **Registry:** Harbor (private projects, quotas, signatures, CVE scanning)
- **AuthN/AuthZ:** OIDC/SAML SSO, LDAP/AD; per namespace/project RBAC
- **Observability:** Prometheus, log shipping; OpenTelemetry exporters

## Technical Specifications

### Availability and Scale

- Control plane: 3 node HA (API server + etcd) with HAProxy/virtual IP; multi AZ capable.
- Data plane: Multiple node pools (CPU/GPU); autoscaling by demand; quotas per project.

### Performance

- Pod density tuned per node flavor; image cache reuse; scheduler and HPA signals optimized.

### Data Protection

- Scheduled etcd snapshots; configuration backup; optional workload/PV backup integration.

JioCloud

## Use Cases

- **CI/CD-driven deployments**
  A developer pushes to Git; the pipeline builds and deploys to a secure Kubernetes cluster— with auto-scaling and healing built-in.

- **ML, AI, and GPU workloads**
  GPU-driven clusters execute model training while still being observable and governed.

- **Hybrid rollouts and secure connectivity**
  Consistent K8s environments deployed across on-prem and cloud in regulated industries. Microservices deployed with Istio-enabled service mesh and RBAC enforcement.

- **Real-time monitoring**
  Built-in dashboards track cluster metrics, pod usage, and alerting across clusters to provide proactive actions.

## Architecture Diagram