

# Managed Kubeflow

## Overview

JioCloud Managed Kubeflow brings the power of open-source Kubeflow into a fully managed, enterprise-grade platform. Your teams stay focused on building value-driven models, while JioCloud handles the complexity of infrastructure, security, and scaling. Instead of juggling fragmented tools, you get one integrated platform for the entire ML journey—offering collaborative notebooks with experiment tracking, reproducible ML pipelines with orchestration, hyperparameter tuning for optimized models, scalable model serving for production, integrated monitoring for reliability, and namespace-level isolation for secure collaboration. No Kubernetes expertise required. No manual patching. No infrastructure headaches. Just a streamlined environment built for enterprise ML at scale.

## Key Features

- Simplified model deployment**  
Expose inference endpoints securely through an API Gateway—no custom serving infrastructure needed.
- Enterprise-grade security**  
Web Application Firewall (WAF) protection for Kubeflow UI and model APIs keeps workloads compliant and resilient.
- Scalable data management**  
Integrated object storage with versioning and persistence for datasets, pipelines, and models.
- Comprehensive observability**  
Get clear visibility into model performance and system health - with alerts before problems impact business.
- Flexible charging model**  
Start small and scale on demand. A base subscription plus pay-as-you-go ensures efficiency across diverse workloads.



# Benefits

## ● Accelerated time-to-market

- 50% faster model deployment through standardized pipelines.
- 20-40% lower operational costs with optimized resource use.
- Elimination of infrastructure bottlenecks that delay projects.

## ● Enhanced governance and security

- Centralized model registry with version control and lineage.
- Comprehensive audit trails for compliance.
- WAF protection for UI and model APIs.

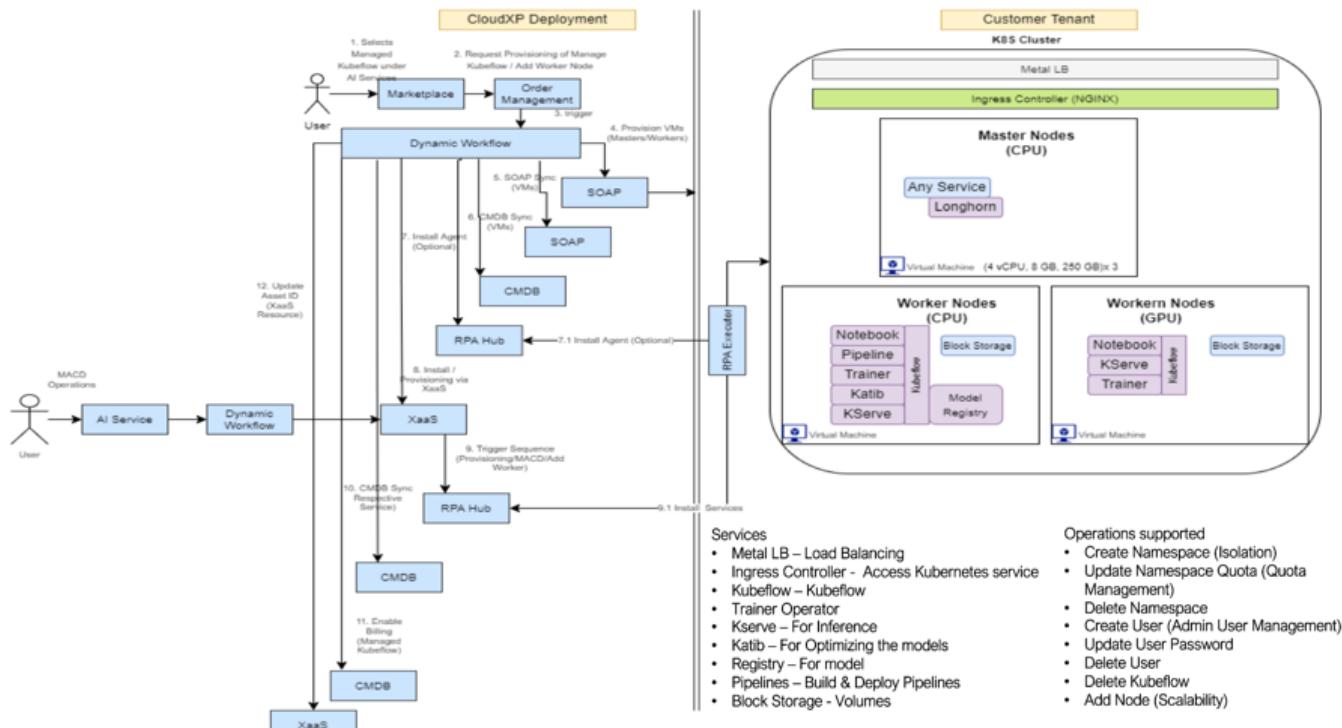
## ● Improved collaboration and scalability

- A shared workspace for data scientists, ML engineers, and DevOps.
- Role-based access for secure, organized collaboration.
- On-demand scaling with simple node provisioning.
- Cross-team visibility into performance and usage.

## Technical Specifications

Category	Details
Operating Systems	Linux (Ubuntu, CentOS), Containerized workloads
Frameworks	TensorFlow, PyTorch
DevOps and MLOps Tools	Kubernetes, Kubeflow, MLflow (integration ready)
Cloud Integrations	Jio Cloud

## Architecture Diagram



## Use Cases

- Standardize ML operations**  
Replace one-off proof-of-concepts with reproducible pipelines across business units.
- Scale securely**  
Run enterprise ML workloads with compliance-ready controls and namespace isolation.
- Enable true collaboration**  
Unite data scientists, ML engineers, and ops teams in a single platform.