

## Charmed Kubeflow on any Kubernetes

Charmed Kubeflow uses charm operators to deliver the 20+ applications that make up the latest version of upstream Kubeflow, for easy consumption anywhere, from workstations, to on-prem, public cloud and edge.

### Artificial Intelligence: the path to production

Data science is transforming a wide range of industries like fin-tech, telco, oil & gas, manufacturing, retail and mobility by providing a competitive edge to early adopters. However, as artificial intelligence reshapes traditional and new industries, factory-like processes that ensure effective operations are crucial to bring AI to production fast.

Although easy to identify, the challenge of operationalizing AI is complex, as it involves:

- Many steps from raw data to inference serving
- Multiple stakeholders like data engineers, data scientists, system administrators, application developers
- Diverse infrastructures from workstation to hybrid-cloud to devices at the edge

In order to address this complexity to create operations around machine learning (MLOps), we draw on Kubernetes.

### Kubeflow, the MLOps toolkit on K8s



The number of diverse tasks involved in creating market-ready AI/ML solutions has proven to be ineffectively handled by sets of notebooks and scripts. These become hard to maintain, share and collaborate on, which leads to technical debt [1] and a long time to

model deployment. To tackle this problem, wider software delivery and ops techniques are adapted to the machine learning (ML) space.

Kubernetes (K8s) is the industry standard for software delivery and operations at scale. Kubeflow provides the cloud-native interface between K8s and data science tools: libraries, frameworks, pipelines and notebooks.


To empower effective team collaboration, Kubeflow hosts a multi-user dashboard for data scientists to access all the tools they need, including features like containerised pipelines, experiment tracking and hyperparameter tuning, while the ops team takes care of the backend.

Visit [ubuntu.com/kubeflow/what-is-kubeflow](https://ubuntu.com/kubeflow/what-is-kubeflow) to learn more.



### Charms for Kubeflow on-rails

Charms are open source universal operators; python code that encapsulates a single app and the automation necessary to operate it, such as how to install and upgrade or how to interact with other applications. Visit [charmhub.io](https://charmhub.io) for the charm collection and [juju.is](https://juju.is) to learn more about operators and lifecycle management.



As Kubeflow is not charmed as a monolithic application, but rather as composable modules, the end user can opt to deploy the full upstream Kubeflow bundle, or customise the deployment to specific needs, e.g. a model development-focused, lightweight desktop version for data scientists.

Visit [charmed-kubeflow.io](https://charmed-kubeflow.io) to know more.

## Easy to get started

Canonical provides a full set of enterprise services to get you started on Kubeflow, from pre-deployment to day-N. These include on-site training, deployment, enterprise-grade support and managed services.

Alternatively, if you want to start small, while evaluating Kubeflow against other technologies, try out Kubeflow on MicroK8s.

Visit [ubuntu.com/kubeflow/install](https://ubuntu.com/kubeflow/install) for a quick trial.

## Multi-cloud portability

Many businesses today choose to operate in hybrid-cloud or multi-cloud scenarios, enjoying the lower-cost compute of on-prem and the elasticity of public clouds.

Thanks to Ubuntu, Charmed Kubernetes and Charmed Kubeflow provide portability of ML workloads across infrastructures, from the datacenter to the public cloud.

## GPU acceleration

Detect and configure GPUs automatically [on MicroK8s](#) and [on Charmed Kubernetes](#) for high-throughput training and inference.

Accelerate the ML workloads on your Kubeflow pipeline with GPU passthrough from machine to Kubernetes to Kubeflow.

## Kubernetes agnostic

Charmed Kubeflow is compatible with any conformant Kubernetes, including AKS, EKS, GKE, MicroK8s, Charmed Kubernetes and any kubeadm-deployed cluster.

Charmed Kubeflow tightly integrates with Canonical's Kubernetes solutions, [MicroK8s](#) for a zero-ops full K8s experience, with automatic high availability, and [Charmed Kubernetes](#) for a configurable deployment at scale.

## Data Lake integration

Cassandra, Hadoop, Spark, Kafka, Hive, PostgreSQL, Elastic. It is common for businesses to have data stacks in production prior to considering Kubeflow adoption.

Canonical stands perfectly positioned to integrate Kubeflow with new, or existing data lakes. Contact us for a free architectural overview.

## Workstation and edge

The composable nature of Charmed Kubeflow allows users to deploy either a pre-specified Kubeflow bundle or a custom set of charms and integrations.

Canonical has created 3 bundles that users can deploy out-of-the-box: full kubeflow; kubeflow-lite for local workstation deployment; and kubeflow-edge inference and distributed training at the edge.

## Canonical Kubeflow services

Charmed Kubeflow provides a lifecycle management layer on top of upstream Kubeflow, consistent with latest versions, so you get full control of your Kubeflow deployment while off-loading complexity to Canonical.

Canonical provides a full set of enterprise services, from evaluation to day-2 operations, which include on-site training, deployment, enterprise-grade support - Ubuntu Advantage - and fully-managed Kubeflow - [Managed Apps](#).

	Training	Deployment	Support	Managed
Service description	On-site 4-day training (up to 15 attendees)	On any conformant Kubernetes	24/7 phone and ticket support	Fully managed Kubeflow (deployment required)
Price	\$29.5k	\$50k	Request quote	Request quote

## Contact us

For more information about Charmed Kubeflow, [contact us](#) or all direct (US Central) +1 737 204 0291 or (UK) +44 203 656 5291