

Charmed Kubeflow for production-grade AI/ML projects

From development to deployment with secure, open-source MLOps tooling



Charmed Kubeflow is an open-source, end-to-end MLOps platform on Kubernetes. It automates machine learning workflows to create a reliable application layer for model development, iteration and production deployment. It is a secure, easy-to-use solution that runs on any environment, from public to private cloud, hybrid or multi-cloud scenarios.

Why Charmed Kubeflow?

Easy to get started

Canonical provides a full set of enterprise services to get you started on Kubeflow, from pre-deployment to day-N. Alternatively, if you want to start small, while evaluating Kubeflow against other technologies, try out Kubeflow on [AWS Marketplace](#).

Secure machine learning tooling

Canonical offers security patching for a wide range of images, including the ones that are specific to Charmed Kubeflow. It protects all the artefacts that are part of the machine learning workflow: from data to models.

Furthermore, Charmed Kubeflow has multi-tenancy capabilities, enabling big teams to work in a secure, portable, reproducible manner.

Kubernetes agnostic

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

Managed Kubeflow

Managed Kubeflow streamlines the operational service delivery of applications by offloading design, implementation and management of Charmed Kubeflow MLOps service environment.

It delivers major gains in efficiency, flexibility and agility for enterprises requiring tailored, optimised service delivery at scale.

A growing MLOps ecosystem

Charmed Kubeflow is the glue of a growing open source MLOps ecosystem and is integrated with a wide range of tooling, including:



Are you looking to learn more about machine learning operations (MLOps)? [Read our guide](#)

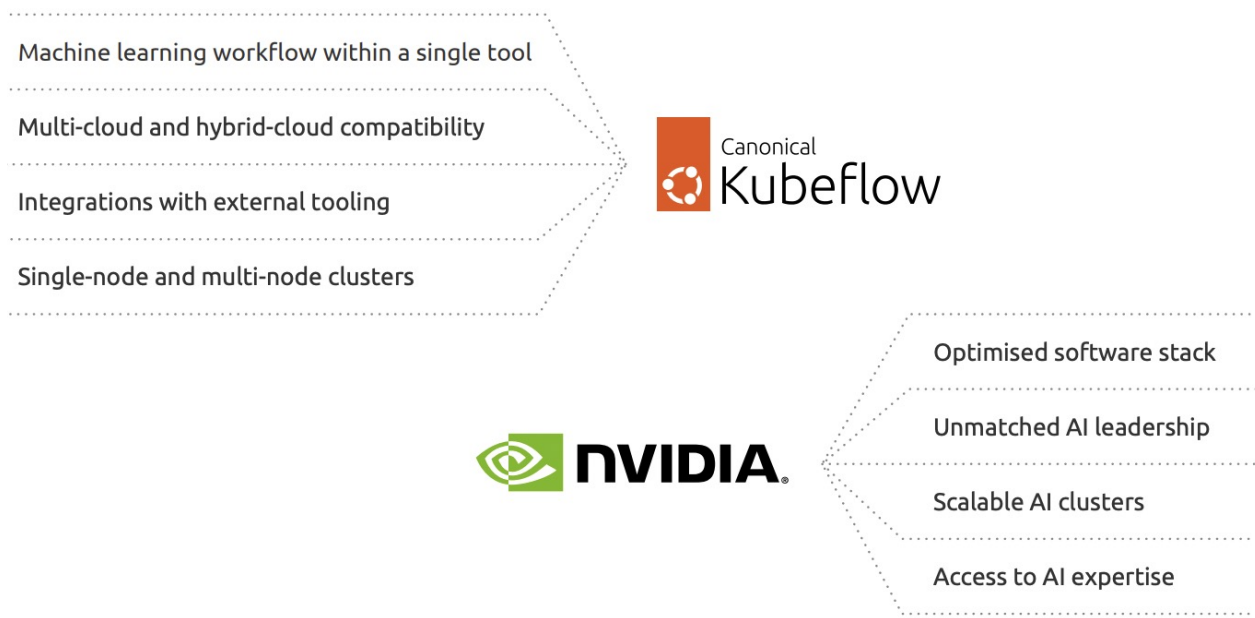
Canonical and NVIDIA: working together to help enterprises benefit from AI/ML

Charmed Kubeflow is certified as part of the DGX-Ready Software program and has been tested on both single-node and multi-node clusters. With NVIDIA AI Enterprise and NVIDIA DGX, Charmed Kubeflow improves the performance of AI workflows, by utilising the hardware to its maximum extent and accelerating project delivery. Charmed Kubeflow can significantly speed up the training phase of AI projects, especially when coupled with DGX systems. With a wide range of integrations, Charmed Kubeflow and NVIDIA DGX bring the MLOps practice to a new level, making the scaling process easier and seamless.

“NVIDIA DGX systems make it possible for every enterprise to access turnkey AI supercomputing to address their most complex AI challenges.”

“DGX-Ready MLOps software like Charmed Kubeflow maximises both the utilisation and efficiency of DGX systems to accelerate the development of AI models.”

— Tony Paikeday, Senior Director of AI Systems at NVIDIA



Canonical services

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

	Consultancy services	Deployment	Support	Managed
Service description	5 types of MLOps advisory services	On any conformant Kubernetes	24/7 phone and ticket support	Fully managed Kubeflow (deployment required)
Price	Read more	\$50k	Contact us	Contact us