

# Charmed Kubeflow delivered on NVIDIA DGX systems

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

Canonical is the company behind Ubuntu, the most popular operating system for container, cloud, and hyperscale computing. NVIDIA leverages Ubuntu as the basis for the NVIDIA DGX™ base operating system, NGC containers and more. With a focus on AI, Canonical's Charmed Kubeflow is certified as part of the DGX-Ready Software program and has been tested on both single-node and multi-node clusters.

Charmed Kubeflow can be run on Canonical Kubernetes and MicroK8s, which provide a hardened, conformant, multi-cloud Kubernetes with full lifecycle automation. Just like Charmed Kubeflow, [Canonical Kubernetes and MicroK8s are certified as part of the DGX-Ready Software program](#), contributing to a seamless experience for the AI practitioner.



# Charmed Kubeflow

## A secure MLOps platform to take machine learning models to production

Canonical's [Charmed Kubeflow](#) provides an open-source, secure, end-to-end MLOps platform that automates machine learning workflows. It runs on any Kubernetes and on any cloud, being the base of an open-source MLOps ecosystem. Charmed Kubeflow is running on NVIDIA DGX™-optimised Ubuntu host images and provides unparalleled integrations and operations, on any compute environment. Combined with NVIDIA AI Enterprise, an end-to-end suite of AI software included with DGX systems, Charmed Kubeflow and DGX provide the tooling AI experts need to train models at scale.

[Charmed Kubeflow](#) translates steps in the data science workflow into Kubernetes jobs. It is one of the official distributions of the [Kubeflow](#) upstream project. Charmed Kubeflow can be used for different purposes, from experimentation using Notebooks to training using the Kubeflow Pipelines or tuning using Katib. Its main goal is to automate ML workloads, encouraging reproducibility and allowing bigger teams to work efficiently.

- Quick deployment
- Run the entire ML lifecycle
- Composable architectures
- Reproducibility, portability, scalability

- Machine learning workflow within a single tool
- Multi-cloud and hybrid-cloud compatibility
- Integrations with external tooling
- Single-node and multi-node clusters



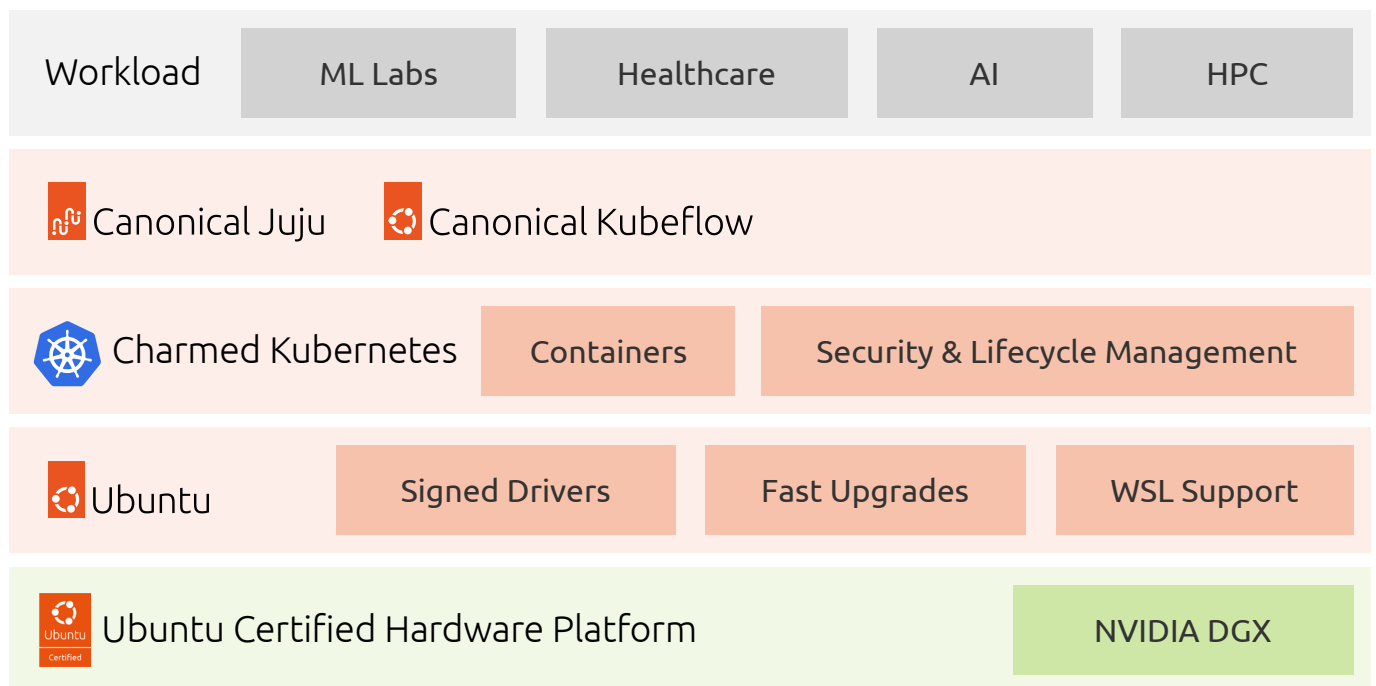
- Optimised software stack
- Unmatched AI leadership
- Scalable AI clusters
- Access to AI expertise

# Advantages of Charmed Kubeflow on NVIDIA DGX

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.

Charmed Kubeflow runs on top of any Kubernetes. However, having it deployed on Canonical's Charmed Kubernetes and MicroK8s, offers compelling integrations and additional features such as resource scheduling. Both products are certified as part of the DGX-Ready Software program as well. With model-driven operations, single-node and high-availability multi-node, you get a complete stack for AI projects that consists of NVIDIA DGX systems, Ubuntu OS, Charmed Kubernetes and Charmed Kubeflow.

- From experimentation to production using one tool
- Security patching, upgrades and updates
- User management
- Runs on supported versions of Kubernetes



## Enterprise security and support

- Day-0 support on the latest upstream Kubeflow versions
- Security maintenance, enterprise support and managed Kubeflow
- Tested on AI-specialised hardware, up-to-date with the latest drivers
- Runs on any Kubernetes
- Cluster lifecycle automation with Juju Open Lifecycle Management
- Enterprise training and workshops with Canonical's [AI/ML advisory](#)

Read more about the [benefits of using compatible hardware and software for AI](#).  
Get started with Charmed Kubeflow on DGX systems.

[Contact us](#) to speed up your AI journey with NVIDIA and Canonical.

For more information about Charmed Kubeflow, please visit [charmed-kubeflow.io/](https://charmed-kubeflow.io/)

To learn about the NVIDIA DGX-Ready Software program, please visit: [nvidia.com/dgx-mlops](https://nvidia.com/dgx-mlops)

To explore more about Kubeflow on DGX systems, please check the [blog](#).