

Reference Architecture for Kubeflow on OpenShift

Accelerate ML/DL workloads using Kubeflow and PowerEdge servers

Customer Results

2 hours vs.
9 months

to run analysis¹

218% ROI

over 3 years²

20 million

images used to train a deep neural network³

Artificial intelligence (AI) and its subsets, machine learning (ML) and deep learning (DL), are gaining widespread adoption for use cases such as computer vision, speech recognition and natural language processing (NLP). Integrating production-grade AI technologies in well-defined platforms within the protection of the data center can facilitate wider adoption of advanced computing, extending investments by supporting AI use cases as well as augmenting the resources available to data science teams.

But not every IT organization has the time and resources to research, integrate and test all the components required to deploy a customized system to run AI workloads. Building a production-ready AI system involves combining various components, often from different vendors, and integrating and managing these disparate pieces. As well, deployments are often tied to a specific cluster, so that moving models — for example, between laptops and cloud clusters or from DevOps to production and back — significantly increases complexity and the chance for human errors.

[Kubeflow](#) together with the Red Hat® OpenShift® Container Platform help address these challenges. Kubeflow is an open-source Kubernetes®-native platform designed to accelerate ML workloads. It's a composable, scalable, portable stack that includes components and automation features to integrate ML tools, so they work together to create a cohesive pipeline that makes it easy to deploy ML applications at scale. Kubeflow requires a Kubernetes environment, such as the Red Hat OpenShift Container Platform, a secure enterprise implementation of open-source Kubernetes.

Running Kubeflow on OpenShift offers several advantages in an ML/DL context. Using Kubernetes as the underlying platform makes models portable, so ML/DL engineers can develop models locally, using a development system such as a laptop, and easily deploy the application to a production Kubernetes environment. In addition, the ability to run ML/DL workloads in the same environment as the rest of your enterprise applications increases control and reduces complexity for IT teams.

Together, Dell Technologies and Red Hat take the guesswork and risk out of AI platform deployment and operations. The Dell Technologies engineering-validated design for AI — OpenShift Container Platform delivers tested, validated, and documented design guidance to help you rapidly deploy Kubeflow and OpenShift on Dell EMC infrastructure.

¹ Dell EMC Case Study, Caterpillar Autonomous Mining, August 2017.

² Forrester Study commissioned by Dell EMC, [The Total Economic Impact of Dell EMC Ready Solutions for AI, Machine Learning with Hadoop](#), August 2018.

³ Dell EMC Video Case Study, [AI startup ZIFF.ai revs up its business with Dell EMC](#), June 2018.

Learn more

- [Streamline your Machine Learning Projects on OpenShift Container Platform v3.11 using KubeFlow](#)
- Learn how to accelerate your ML/DL projects using KubeFlow and NVIDIA GPUs: [Executing ML/DL Workloads using OpenShift Container Platform v3.11 white paper](#)
- Explore the KubeFlow project: [KubeFlow: The Machine Learning Toolkit for Kubernetes](#)
- Contact the Dell Technologies OpenShift team: openshift@dell.com

Dell Technologies reference architectures

Dell Technologies and Red Hat offer an engineering-tested and proven design for ML/DL workloads using enterprise-grade Kubernetes container orchestration. This enterprise-ready platform serves as the foundation for building a robust, high-performance environment that supports various lifecycle stages of an AI project: model development using Jupyter® Notebooks, rapid iteration and testing using Tensorflow™, training DL models using GPUs, and enabling prediction using developed models.

The [Dell EMC Running ML/DL Workloads Using Red Hat OpenShift Container Platform white paper](#) outlines the steps for building a powerful AI platform so that compute, memory, storage and networking resources are effectively and efficiently utilized to accelerate ML/DL workloads. It describes how to deploy KubeFlow on OpenShift Container Platform using Dell EMC PowerEdge servers with NVIDIA Tesla® GPUs to achieve a high-performance AI environment for ML/DL scientists without having to build a complete platform from scratch.

Expert Dell Technologies advanced computing teams are ready to help you create ML/DL configurations from our portfolio of workstations, servers, storage, networking, software and services. Dell Technologies expertise is enhanced through collaboration with Red Hat engineers, worldwide Dell Technologies [Customer Solution Centers](#), [HPC & AI Centers of Excellence](#), the [HPC & AI Innovation Lab](#) and the broader AI community. These experts can work with you to create a solution that balances price and performance for your requirements. With an expansive portfolio of infrastructure and AI-enabled IT solutions, including dedicated AI consulting services, Dell Technologies solutions allow you to start small and grow according to your own AI journey.

Red Hat and Dell Technologies

Red Hat is the world's leading provider of enterprise open-source solutions, using a community-powered approach to deliver high-performing Linux®, cloud, container and Kubernetes technologies. Red Hat can help you standardize across environments, develop cloud-native applications, and integrate, automate, secure and manage complex environments.

Dell Technologies enables organizations to modernize, automate and transform their data center using industry-leading converged infrastructure, servers, storage and data protection technologies. Businesses get a trusted foundation to transform their IT and develop new and better ways to work through hybrid cloud, the creation of cloud-native applications, and big data solutions.

