



There's no single way to build and operationalize ML models, but there is a consistent need to gather and prepare data, develop models, turn models into intelligent applications, and derive revenue from those applications. Adopting MLOps practices means there's no time wasted building or deploying a model and keeping it up to date. Red Hat® OpenShift®, a leading hybrid cloud application platform powered by Kubernetes, in combination with Red Hat OpenShift AI, includes key capabilities to enable MLOps in a consistent way across datacenters, public cloud computing, and edge computing.

## MLOps: Machine learning operations with Red Hat OpenShift



Red Hat OpenShift Data Science extends OpenShift to build, serve, and monitor ML models. Part of Red Hat OpenShift AI, Red Hat OpenShift Data Science provides an open modular data science platform where data scientists can use their preferred tools or draw from an ecosystem of partner solutions to generate insights from data. In addition to common model development tools like Jupyter and associated frameworks, OpenShift Data Science provides serving, data pipelines, and monitoring tooling to help deploy models more quickly into production.

## Enhance MLOps with Red Hat OpenShift Data Science





Red Hat OpenShift AI is a portfolio of products, including OpenShift Data Science, that enables a consistent and scalable infrastructure across the full lifecycle of AI/ML to achieve production-ready AI models and applications

## 1 Model training

ML models are trained on Red Hat OpenShift Data Science, running on Red Hat OpenShift. Alternatively, you could use the open source Jupyter notebook technology on OpenShift to build models.

## 2 Automate model training

Red Hat OpenShift Pipelines, fully integrated with OpenShift Data Science, is an event-driven, continuous integration capability that helps create pipelines to train machine learning (ML) models using secure container images:

- ▶ Clearly setting the model training steps, using secured and immutable container images.
- ▶ Saving the models ready for deployment in a model store or including the saved models to container images with Red Hat OpenShift build.
- ▶ Testing the containerized model images to ensure they remain functional.

## 3 More securely deploy models anywhere

Declarative configuration managed by Red Hat OpenShift GitOps automates the deployment of ML models at scale, anywhere, by:

- ▶ **Configuring** Red Hat OpenShift environments for artificial intelligence (AI) inferencing, anywhere, via Git repositories. These configuration requirements are recorded and can be versioned at the source, reducing the propensity for errors and increasing developer and data scientist productivity.

- ▶ **Monitoring** the manifest with the latest model version to be used by the intelligent app. OpenShift Data Science model server is always updated for more secure deployments.
- ▶ **Triggering** the deployment of the latest version of the models as ML services that are used by intelligent software applications via application programming interfaces (APIs).
- ▶ **Deploying** the latest versions of the ML models and the associated intelligent application to Red Hat OpenShift at all the locations where ML inferencing is to be performed (i.e., datacenters, public cloud computing, and edge computing).

## 4 Monitor model drift

Monitor models for reliability, speed, scale, etc. with the tooling from OpenShift Data Science or one of our ecosystem partners, and update with retraining and redeployment, as needed.

### Learn more

#### MLOps checklist:

Top 5 ways to implement MLOps successfully ▶

#### Red Hat OpenShift Data Science:

Accelerate your data science ▶

#### AI/ML environments e-book:

Build a production-ready AI/ML environment ▶