

MLOps Practices with Red Hat OpenShift AI (AI500)

Duration: 5 Days **Course Code:** AI500

Overview:

Experience the possibilities of MLOps through proven open culture and practices used by Red Hat to support customer innovation.

MLOps Practices with Red Hat OpenShift AI (AI500) is a five-day immersive class, offering attendees an opportunity to experience and implement a successful MLOps adoption journey. While many AI or data science training programs focus on a particular framework or technology, this course covers how the best Open Source tools fit together in a full MLOps workflow. It blends continuous discovery, continuous training, and continuous delivery in a highly engaging experience simulating real-world machine learning scenarios.

To achieve the learning objectives, participants should include multiple roles from across the organization. Data scientists, machine learning engineers, platform engineers, architects, and product owners will gain experience working beyond their traditional silos.

The daily routine simulates a real-world delivery team, where cross-functional teams learn how collaboration breeds innovation.

Armed with shared experiences and best practices, the team can apply what it has learned to help the organization's culture and mission succeed in the pursuit of new projects and improved processes.

This course introduces real-world MLOps culture principles and modern practices. You will develop a predictive machine learning model using Red Hat OpenShift and Red Hat OpenShift AI, and other industry-standard MLOps software, tools, and techniques.

By the end of the course, you will be equipped to apply MLOps principles and leverage open-source solutions to drive and lead AI transformation initiatives within your organization.

This course is based on Red Hat OpenShift AI, Red Hat OpenShift GitOps and Predictive AI

Target Audience:

This experience demonstrates how individuals across different roles must learn to share, collaborate, and work toward a common goal to achieve positive outcomes and drive innovation.

It is especially valuable for: MLOps Platform Users: Data scientists, data engineers, and application developers. MLOps Platform Providers:

Machine learning engineers, MLOps engineers, and platform engineers. MLOps Platform Stakeholders: Architects and IT managers.

The scenario incorporates technical aspects of working with machine learning systems, offering practical insights into how these roles can align their efforts.

You will learn how to continuously deliver value to your customers by accelerating the deployment of new models to market. Our instructors will share experiences and best practices learned from engaging directly with customers during Red Hat services engagements.

Note: Starting January 2026 this course only exists in CR (classroom) if scheduled or Closed course modalities. No RHLS-Course for this course.

Updated Jan2026

Objectives:

- This course takes you an end to end journey of a Predictive Intelligent Application use case, from ideation to inner loop experimentation to production, while bringing different personas together to seamlessly collaborate on a single platform.
- After this course participants should be able to:
- Experience MLOps culture, explore MLOps practices, and apply learning to bring a machine learning model into production
- Apply MLOps principles to streamline the development and deployment of machine learning models
- Gain hands-on experience with modern tools and processes, covering the entire lifecycle from inner loop development to outer loop operations
- Enhance skills in collaborative coding styles with pair and mob programming style

Prerequisites:

- Containers, Kubernetes and Red Hat OpenShift Technical Overview (DO080) or Basic understanding of OpenShift/Kubernetes and containers is helpful
- High level understanding of AI or Red Hat AI Foundations is beneficial

Take Red Hat free assessment to gauge whether this offering is the best fit for your skills [Red Hat Skills Assessment](#)

Testing and Certification

None

Follow-on-Courses:

None

Content:

What is MLOps?	Introduction to MLOps: a set of practices that automate and simplify machine learning workflows and deployments.	Properly handle pre- and post-processing for data and predictions, explore autoscaling to handle loads, and introduce advanced deployment patterns like canary and blue-green deployments to ensure safe and seamless model rollouts.
Brainstorm and explore what principles, practices, and cultural elements make up a MLOps model for ML model developments and deployments.	Here we will create our MLOps environment where the continuous training pipeline, automated deployment, and the supporting toolings will be running.	Feature Stores
Inner Loop	Monitoring	Robust ways of dealing with data features and their changes, as well as making sure features are homogeneous between training and serving.
Familiarize ourselves with the necessary tools for experimenting and building our model; we will create a workbench, explore the dataset, start tracking our experiments, and deploy our models.	Machine learning models can be influenced by various factors, including changes in data patterns, shifts in user behavior, and evolving external conditions. By implementing continuous monitoring, we will proactively identify these changes, assess their impact on model accuracy, and make necessary adjustments to maintain optimal performance.	Security
Training Pipelines	Data Versioning	Implement automated security guardrails to stay compliant with the organizations security practices and extend them to the models.
Transition to automating the previous steps for productionizing our model training.	Enhance traceability by introducing versioning for our datasets as they change over time.	
Outer Loop	Advanced Deployments	

Additional Information:

Official course book provided to participants

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/