

MLOps Practices with Red Hat OpenShift AI

Course code: AI500

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

Affiliate	Duration	Course price	ITB
Praha	5	5 200 €	0
Brno	5	5 200 €	0
Bratislava	5	5 200 €	0

The prices are without VAT.

Course terms

Date	Duration	Course price	Type	Course language	Location
------	----------	--------------	------	-----------------	----------

The prices are without VAT.

Who is the course for

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.

What we teach you

- 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.
- This course blends cultural and technical practices into a unique, highly-engaging experience, packed with real-world applications. You will learn MLOps practices and how they build upon each other to improve team alignment and delivery efficiency.
- Most AI training focuses on a particular framework or technology, this course combines the best Open Source tools while giving you the experience of how they fit together to reliably and efficiently build, deploy and maintain AI models in production.

GOPAS Praha
Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno
Nové sady 996/25
602 00 Brno
Tel.: +420 530 513 590
info@gopas.cz

GOPAS Bratislava
Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk



Copyright © 2026 GOPAS, a.s.,
All rights reserved

MLOps Practices with Red Hat OpenShift AI

Required skills

- Take our free assessment
- to gauge whether this offering is the best fit for your skills.
- Containers, Kubernetes and Red Hat OpenShift Technical Overview (D0080)
- or Basic understanding of OpenShift/Kubernetes and containers is helpful
- High level understanding of AI or
- Red Hat AI Foundations
- is beneficial

Course outline

What is MLOps?

- Brainstorm and explore what principles, practices, and cultural elements make up a MLOps model for ML model developments and deployments.

Inner Loop

- 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.

Training Pipelines

- Transition to automating the previous steps for productionizing our model training.

Outer Loop

- Introduction to MLOps: a set of practices that automate and simplify machine learning workflows and deployments.
- Here we will create our MLOps environment where the continuous training pipeline, automated deployment, and the supporting toolings will be running.

Monitoring

- 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.

Data Versioning

- Enhance traceability by introducing versioning for our datasets as they change over time.

Advanced 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.

Feature Stores

- Robust ways of dealing with data features and their changes, as well as making sure features are homogeneous between training and serving.

Security

- Implement automated security guardrails to stay compliant with the organizations security practices and extend them to the models.