

Developing and Deploying AI/ML Applications on Red Hat OpenShift AI

Course code: AI267

Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267) provides students with the fundamental knowledge about using Red Hat OpenShift for developing and deploying AI/ML applications. This course helps students build core skills for using Red Hat OpenShift AI to train, develop and deploy machine learning models through hands-on experience. This course is based on Red Hat OpenShift @ 4.14, and Red Hat OpenShift AI 2.8. Note: This course is offered as a 3 day in person class, a 4 day virtual class or is self-paced. Durations may vary based on the delivery. For full course details, scheduling, and pricing, select your location then "get started" on the right hand menu.

Pro koho je kurz určen

- Data scientists and AI practitioners who want to use Red Hat OpenShift AI to build and train ML models
- Developers who want to build and integrate AI/ML enabled applications
- MLOps engineers responsible for installing, configuring, deploying, and monitoring AI/ML applications on Red Hat OpenShift AI

Co Vás naučíme

Impact on the Organization

Organizations collect and store vast amounts of information from multiple sources. With Red Hat OpenShift AI, organizations have a platform ready to analyze data, visualize trends and patterns, and predict future business outcomes by using machine learning and artificial intelligence algorithms.

Impact on the Individual

As a result of attending this course, you will understand the foundations of the Red Hat OpenShift AI architecture. You will be able to install Red Hat OpenShift AI, manage resource allocations, update components and manage users and their permissions. You will also be able to train, deploy and serve models, including how to use Red Hat OpenShift AI to apply best practices in machine learning and data science. Finally you will be able to create, run, manage and troubleshoot data science pipelines.

- Introduction to Red Hat OpenShift AI
- Data Science Projects
- Jupyter Notebooks
- Installing Red Hat OpenShift AI
- Managing Users and Resources
- Custom Notebook Images
- Introduction to Machine Learning
- Training Models
- Enhancing Model Training with RHOAI
- Introduction to Model Serving
- Model Serving in Red Hat OpenShift AI
- Introduction to Workflow Automation
- Elyra Pipelines
- KubeFlow Pipelines

Požadované vstupní znalosti

- Experience with Git is required
- Experience in Python development is required, or completion of the Python Programming with Red Hat (AD141) course
- Experience in Red Hat OpenShift is required, or completion of the Red Hat OpenShift Developer II: Building and Deploying Cloud-native Applications (D0288) course

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

Developing and Deploying AI/ML Applications on Red Hat OpenShift AI

- Basic experience in the AI, data science, and machine learning fields is recommended

Osnova kurzu

Introduction to Red Hat OpenShift AI

- Identify the main features of Red Hat OpenShift AI, and describe the architecture and components of Red Hat AI.

Data Science Projects

- Organize code and configuration by using data science projects, workbenches, and data connections

Jupyter Notebooks

- Use Jupyter notebooks to execute and test code interactively

Installing Red Hat OpenShift AI

- Installing Red Hat OpenShift AI by using the web console and the CLI, and managing Red Hat OpenShift AI components

Managing Users and Resources

- Managing Red Hat OpenShift AI users, and resource allocation for Workbenches

Custom Notebook Images

- Creating custom notebook images, and importing a custom notebook through the Red Hat OpenShift AI dashboard

Introduction to Machine Learning

- Describe basic machine learning concepts, different types of machine learning, and machine learning workflows

Training Models

- Train models by using default and custom workbenches

Enhancing Model Training with RHOAI

- Use RHOAI to apply best practices in machine learning and data science

Introduction to Model Serving

- Describe the concepts and components required to export, share and serve trained machine learning models

Model Serving in Red Hat OpenShift AI

- Serve trained machine learning models with OpenShift AI

Custom Model Servers

- Deploy and serve machine learning models by using custom model serving runtimes

Introduction to Data Science Pipelines

- Create, run, manage, and troubleshoot data science pipelines

Elyra Pipelines

- Creating a Data Science Pipeline with Elyra

KubeFlow Pipelines

- Creating a Data Science Pipeline with KubeFlow SDK