

AI267 Developing and Deploying AI/ML Applications on Red Hat OpenShift

Course Description

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 ® 4.16, and Red Hat OpenShift AI 2.13.

Course content summary

- Introduction to Red Hat OpenShift AI
- Data Science Projects
- Jupyter Notebooks
- Red Hat OpenShift AI Installation
- Users and Resources Management
- 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 Data Science Pipelines
- Working with Pipelines
- Controlling Pipelines and Experiments

Audience for this course

- 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
- Developers, data scientists, and AI practitioners who want to automate their ML workflows
- MLOps engineers responsible for operationalizing the ML lifecycle on Red Hat OpenShift AI

Recommended training

- Experience with Git is required
- Experience in Python development is required, or completion of the <u>Python Programming with Red</u> <u>Hat (AD141)</u> course
- Experience in Red Hat OpenShift is required, or completion of the <u>Red Hat OpenShift Developer II:</u> <u>Building and Deploying Cloud-native Applications (DO288)</u> course
- Basic experience in the AI, data science, and machine learning fields is recommended

Outline for this course

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

Red Hat OpenShift AI Installation

Install Red Hat OpenShift AI and manage Red Hat OpenShift AI components

User and Resource Management

Manage Red Hat OpenShift AI users and allocate resources

Custom Notebook Images

Create and import custom notebook images in Red Hat OpenShift AI

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 Al

Serve trained machine learning models with OpenShift AI

Introduction to Data Science Pipelines

Define and set up Data Science Pipelines

Working with Pipelines

Create data science pipelines with the Kubeflow SDK and Elyra

Controlling Pipelines and Experiments

Configure, monitor, and track pipelines with artifacts, metrics, and experiments

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 define and set up data science pipelines with Red Hat OpenShift AI.