

DEVELOPING AND DEPLOYING AI/ML APPLICATIONS ON RED HAT OPENSIFT AI WITH EXAM (AI268)

Course Code: 832011

An introduction to developing and deploying AI/ML applications on Red Hat OpenShift AI.

An introduction to developing and deploying AI/ML applications on Red Hat OpenShift AI.

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. The Red Hat Certified Specialist in OpenShift AI Exam (EX267) is included in the offering.

Related offer

Get the Developing and Deploying AI/ML Applications on Red Hat OpenShift AI with Exam (AI268) at 15% off local MSRP, now through September 30, 2024. You must select both the course plus the exam to receive this offer. Discount will be applied to the local MSRP of eligible offerings at checkout. This offer may not be combined with other offers or discounts.

Course Content Summary

- 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

What You'll Learn

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.

Who Needs to Attend

- 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

Prerequisites

- 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 (DO288) course
- Basic experience in the AI, data science, and machine learning fields is recommended

Visit us at www.globalknowledge.com or call us at 1-866-716-6688.

Date created: 5/26/2026 12:44:48 AM

Copyright © 2026 Global Knowledge Training LLC. All Rights Reserved.