Red Hat Certified Specialist in OpenShift AI (EX267)

ID EX267 **Price** 477.— €(excl. tax) **Duration**

Course Overview

The Red Hat Certified Specialist in OpenShift AI exam tests candidates' ability to deploy OpenShift AI and configure it to build, deploy, and manage machine learning models to support AI-enabled applications.

By passing this exam, you become a Red Hat Certified Specialist in OpenShift AI that also counts towards earning a Red Hat Certified Architect (RHCA®).

This exam is based on Red Hat OpenShift AI version 2.25 and Red Hat OpenShift Container Platform version 4.18.

Who should attend

- System and Software Architects who want to validate the ability to design and integrate scalable AI/ML infrastructure using Red Hat OpenShift AI
- Developers who want to demonstrate proficiency in implementing and automating MLOps workflows and integrating models into production applications
- Data Scientists who want to prove expertise in developing, training, serving, and monitoring models within Red Hat OpenShift AI

Prerequisites

Candidates for this exam should:

- Have taken Red Hat OpenShift Developer II: Building and Deploying Cloud-native Applications (DO288) or have comparable work experience using OpenShift Container Platform
- Have taken Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267) or have comparable work experience using the features of OpenShift AI
- Review the Red Hat Certified Specialist in OpenShift AI exam (EX267) objectives

Take our free assessment to find the course that best supports your preparation for this exam.

Preparation

Red Hat encourages you to consider taking the course Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267) to help prepare. Attendance in these classes is not required; students can choose to take just the exam.

While attending Red Hat classes can be an important part of your preparation, attending class does not guarantee success on the exam. Previous experience, practice, and native aptitude are also important determinants of success.

Many books and other resources on system administration for Red Hat products are available. Red Hat does not endorse any of these materials as preparation guides for exams. Nevertheless, you may find additional reading helpful to deepen your understanding.

Course Content

Candidates for the Red Hat Certified Specialist in OpenShift AI should be able to accomplish the following tasks. Relevant product-specific documentation will be provided, but candidates should be prepared to perform these tasks without assistance.

Understand Red Hat OpenShift AI architecture and fundamentals

- Understand Red Hat OpenShift AI's relationship with OpenShift Container Platform
- Understand MLOps, GenAIOps, and AI/ML concepts
- Know how Red Hat OpenShift AI components work in data science projects

Manage data science projects and workbenches

- Create, configure, and manage projects and permissions
- Create and edit workbenches with custom images, versions, and sizes
- Build and import custom workbench images
- Monitor resource usage and training processes with TensorBoard

Configure data connections

- Create connections (S3, database, etc.)
- Store and retrieve data and artifacts from external services

Identify and allocate resources

- Use nodeSelectors and tolerations
- Allocate workbenches and model servers to specific nodes

Deploy and serve models

- Understand model serving workflow and KServe architecture
- Deploy models using Standard and Advanced modes
- Store models in S3 buckets, OCI containers, or PVCs
- Serve predictive models with OpenVINO runtime
- Deploy and serve LLMs with vLLM runtime
- Create and configure custom serving runtimes

Manage models with the Model Registry

- Package models as OCI image artifacts
- Register and version models in the Model Registry
- Deploy models from the Model Registry
- Query the Model Registry API

Monitor AI models and performance

- Monitor model bias and data drift with TrustyAI
- Monitor hardware consumption with OpenShift monitoring stack and Grafana
- Analyze resource utilization and optimize based on monitoring insights

Create and manage data science pipelines

- Create pipeline servers and pipelines with Elyra and KubeFlow SDK
- Use container components and manage artifacts
- Configure Kubernetes features in pipelines
- Use experiments to compare pipeline runs

Optimize and evaluate models

- Select models from Red Hat OpenShift AI catalog and Hugging Face
- Optimize models with LLM Compressor (compression and quantization)
- Evaluate LLM performance with LMEval using standard and custom benchmarks

Build GenAI applications

- Understand and apply GenAI application patterns
- Build simple GenAI applications with streaming responses
- Build RAG applications with vector databases and document processing
- Build agentic applications with tools and multi-step reasoning
- Implement guardrails for content safety and input/output validation

Collaborate with Git and develop ML models

- Manage Jupyter notebooks with Git version control
- Train models in Python using foundational ML libraries
- Load data scalably and save or export models

Deploy and store models

- Deploy models using Red Hat OpenShift AI interface (Standard and Advanced modes)
- Store models using S3 buckets, OCI containers, or persistent volume claims
- Understand supported model storage locations
- Configure model deployment settings

Exam format

This exam is a performance-based evaluation of skills and knowledge required to configure and manage Red Hat OpenShift AI. Candidates perform routine tasks using Red Hat OpenShift Container Platform and Red Hat OpenShift AI and are evaluated on whether they have met specific objective criteria. Performance-based testing means that candidates must perform tasks similar to what they perform on the job.

This exam consists of a single section lasting four hours.

Scores and reporting

Official scores for exams come exclusively from Red Hat Certification Central. Red Hat does not authorize examiners or training partners to report results to candidates directly. Scores on the exam are usually reported within 3 U.S. business days.

Exam results are reported as total scores. Red Hat does not report performance on individual items, nor will it provide additional information upon request.

About Fast Lane

Fast Lane is a global, award-winning specialist in technology and business training as well as consulting services for digital transformation. As the only global partner of the three cloud hyperscalers- Microsoft, AWS and Google- and partner of 30 other leading IT vendors, Fast Lane offers qualification solutions and professional services that can be scaled as needed. More than 4,000 experienced Fast Lane professionals train and advise customers in organizations of all sizes in 90 countries worldwide in the areas of cloud, artificial intelligence, cyber security, software development, wireless and mobility, modern workplace, as well as management and leadership skills, IT and project management.

Fast Lane Services

- ✓ High End Technology Training
- ✓ Business & Soft Skill Training
- ✓ Consulting Services
- ✓ Managed Training Services
- ✓ Digital Learning Solutions
- ✓ Content Development
- ✓ Remote Labs
- ✓ Talent Programs
- ✓ Event Management Services

Training Methods

- ✓ Classroom Training
- ✓ Instructor-Led Online Training
- ✓ FLEX Classroom – Classroom & Online Hybrid
- ✓ Onsite & Customized Training
- ✓ E-Learning
- ✓ Blended & Hybrid Learning
- ✓ Mobile Learning

Technologies & Solutions

- ✓ Digital Transformation
- ✓ Artificial Intelligence
- ✓ Cloud
- ✓ Networking
- ✓ Cyber Security
- ✓ Wireless & Mobility
- ✓ Modern Workplace
- ✓ Data Center

Worldwide Presence
with high-end training centers around the globe

Multiple Awards
from vendors such as AWS, Microsoft, Cisco, Google, NetApp, VMware

Experienced SMEs
with over 19.000 combined certifications

Germany
Fast Lane Institute for Knowledge Transfer GmbH
Tel. +49 40 25334610
info@flane.de / www.flane.de

Austria
ITLS GmbH
(Partner of Fast Lane)
Tel. +43 1 6000 8800
info@itls.at / www.itls.at

Switzerland
Fast Lane Institute for Knowledge Transfer (Switzerland) AG
Tel. +41 44 8325080
info@flane.ch / www.flane.ch