

Red Hat OpenShift AI Technical Overview

Course Duration: 16 Hours

Course code: AI067

1. Course Overview

AI067 – Red Hat OpenShift AI Technical Overview provides a foundational understanding of Red Hat’s AI/ML platform capabilities. This free, on-demand course introduces the OpenShift AI architecture, MLOps workflows, and integration with large language models (LLMs) and image processing tools. Participants will explore how to build, train, and deploy AI solutions efficiently within hybrid cloud environments. Designed for IT professionals, developers, and data scientists, it offers a practical insight into accelerating AI adoption using Red Hat OpenShift AI.

2. What you’ll learn?

By the end of this course, learners will:

This course aims to introduce learners to the core concepts, architecture, and capabilities of Red Hat OpenShift AI. Participants will explore how OpenShift AI supports the entire machine learning lifecycle, from data preparation to model deployment, using container-native tools and scalable infrastructure. By the end of the course, learners will understand how to leverage OpenShift AI to streamline collaboration between data science and operations teams, accelerating time to value for AI projects.

3. Target Audience

- Data scientists and data engineers
- Machine learning engineers
- DevOps professionals supporting AI/ML workflows
- Platform engineers managing container platforms
- IT architects and solution designers
- Professionals interested in deploying AI/ML workloads on OpenShift

4. Pre-Requisites

Participants should have:

- Participants should have a foundational understanding of Linux systems and containers, as well as a basic familiarity with Kubernetes concepts. Prior experience with data science or machine learning workflows will help learners better contextualize how OpenShift AI supports MLOps practices.
- Basic knowledge of Linux operating system
- Familiarity with container technologies (e.g., Docker)
- General awareness of Kubernetes concepts
- Introductory understanding of data science or machine learning processes (recommended)

5. Course content

Module 01 – Introduction to Event-Driven Architecture

- Overview of Event-Driven Applications and Use Cases
- Introduction to Apache Kafka and Red Hat AMQ Streams
- Understanding Core Kafka Concepts: Topics, Partitions, and Brokers

Module 02 – Setting Up the Development Environment

- Installing and Configuring Apache Kafka and AMQ Streams
- Integrating Kafka with Red Hat OpenShift
- Using Kafka Command-Line Tools for Basic Operations

Module 03 – Producing and Consuming Events

- Creating Kafka Producers and Consumers in Java
- Managing Message Serialization and Deserialization
- Implementing Consumer Groups and Offsets

Module 04 – Designing Event Streams

- Best Practices for Topic Design and Partitioning
- Message Ordering and Delivery Guarantees
- Event Schemas and Schema Registry Usage

Module 05 – Stream Processing and Patterns

- Introduction to Kafka Streams API
- Implementing Streaming Patterns (Filter, Aggregate, Join)
- Handling Late or Out-of-Order Events

Module 06 – Reliability, Scalability, and Fault Tolerance

- Configuring Replication and Acknowledgements
- Error Handling and Retry Strategies
- Ensuring High Availability in Kafka Deployments

Module 07 – Monitoring, Security, and Deployment

- Securing Kafka Clusters with Authentication and Authorization
- Monitoring Kafka with Metrics and Logging
- Deploying Event-Driven Applications in Hybrid and Cloud Environments

