

MLOps Engineering on AWS

Course Duration: 24 Hours

Course code:AMWSMLOPS

1. Course Overview

During this 3-day course, you focus on how to plan and deploy machine learning workloads with AWS services, and how to perform operational tasks such as model training, deployment, monitoring, and scaling. Learn to configure logging, alerts, and metrics for ML pipelines. Deploy AWS SageMaker projects, integrate with AWS Lambda, Step Functions, and manage CI/CD for ML. Understand how to set up observability for ML models using Amazon CloudWatch and Amazon SageMaker Model Monitor. The concepts learned in this course should be transferable for engineers who need to plan, deploy, and manage MLOps workflows on AWS.

2. What you'll learn?

By the end of the course, you should be able to meet the following objectives:

- Describe MLOps principles and AWS ML services
- Explain the AWS SageMaker ecosystem and integration with other AWS services
- Plan and deploy end-to-end MLOps pipelines on AWS
- Identify use cases for MLOps on AWS
- Perform operational tasks such as automated model training, deployment, and monitoring
- Configure logging, alerts, and metrics for ML workflows
- Deploy AWS Lambda and Step Functions for ML automation
- Implement model monitoring using Amazon SageMaker Model Monitor
- Integrate CI/CD tools with AWS for ML lifecycle management

3. Target Audience

ML engineers, data scientists, DevOps engineers, and platform operators responsible for deploying and managing ML workloads on AWS.

4. Pre-Requisites

Familiarity with:

- AWS Cloud fundamentals
- Machine Learning workflows
- Basic knowledge of CI/CD concepts
- Python programming

5. Course content

Module 1: Course Introduction

- Introduction
- Course Contents

Module 2: Introduction to MLOps on AWS

- Introduction
- What is MLOps
- How AWS supports MLOps workflows
- When to use AWS MLOps services
- Summary

Module 3: AWS SageMaker Core Components

- SageMaker Studio
- SageMaker Notebooks
- Training Jobs and Endpoints

Module 4: AWS Data Management for MLOps

- Data ingestion with AWS Glue and Amazon S3
- Data preprocessing pipelines with SageMaker Processing
- Configuring data encryption at rest and in transit

Module 5: Networking and Security for MLOps on AWS

- VPC configurations for SageMaker
- PrivateLink and network isolation
- Configuring IAM roles and permissions for ML workloads

Module 6: Planning an MLOps Deployment on AWS

- Overview of MLOps deployment planning
- Designing pipeline architecture for ML lifecycle

Module 7: Deploying MLOps Pipelines with SageMaker

- Deploying training and inference pipelines
- Connecting to on-premise or hybrid ML environments
- Configuring SageMaker endpoints and auto-scaling

Module 8: Operational Commands and Automation in MLOps

- Viewing pipeline execution status
- Cancelling or retrying pipeline runs
- Integrating CI/CD with CodePipeline and CodeBuild

Module 9: Configuring Logging and Monitoring in AWS MLOps

- Setting up Amazon CloudWatch Logs
- Configuring diagnostic settings for ML pipelines

Module 10: Configuring Alerts and Working with Metrics in MLOps

- Supported ML model metric
- Configuring alert rules in CloudWatch
- Working with performance dashboards

Module 11: Deploying AWS Lambda and Step Functions for ML Automation

- Overview of AWS automation for MLOps
- Orchestrating training and deployment with Step Functions
- Automating retraining triggers
- Managing Lambda-based inference APIs
- Updating automation workflows
- Removing outdated automation pipelines

Module 12: Model Monitoring with Amazon SageMaker Model Monitor

- Prerequisites
- Setting up monitoring schedules
- Monitoring data drift and model performance
- Integrating alerts with CloudWatch
- Known limitations