

Module Title : MLOps Engineering on AWS

Duration : 3 days

Description

This course builds upon and extends the DevOps practice prevalent in software development to build, train, and deploy machine learning (ML) models. The course stresses the importance of data, model, and code to successful ML deployments. It will demonstrate the use of tools, automation, processes, and teamwork in addressing the challenges associated with handoffs between data engineers, data scientists, software developers, and operations. The course will also discuss the use of tools and processes to monitor and take action when the model prediction in production starts to drift from agreed-upon key performance indicators.

The instructor will encourage the participants in this course to build an MLOps action plan for their organization through daily reflection of lesson and lab content, and through conversations with peers and instructors

- Course level: Intermediate
- Duration: 3 days

Intended Audience

This course is intended for any one of the following roles with responsibility for productionizing machine learning models in the AWS Cloud:

- DevOps engineers
- ML engineers
- Developers/operations with responsibility for operationalizing ML models

Course Objectives

In this course, you will learn to:

- Describe machine learning operations
- Understand the key differences between DevOps and MLOps
- Describe the machine learning workflow
- Discuss the importance of communications in MLOps
- Explain end-to-end options for automation of ML workflows
- List key Amazon SageMaker features for MLOps automation
- Build an automated ML process that builds, trains, tests, and deploys models
- Build an automated ML process that retrains the model based on change(s) to the model code
- Identify elements and important steps in the deployment process

- Describe items that might be included in a model package, and their use in training or inference
- Recognize Amazon SageMaker options for selecting models for deployment, including support for ML frameworks and built-in algorithms or bring-your-own-models
- Differentiate scaling in machine learning from scaling in other applications
- Determine when to use different approaches to inference
- Discuss deployment strategies, benefits, challenges, and typical use cases
- Describe the challenges when deploying machine learning to edge devices
- Recognize important Amazon SageMaker features that are relevant to deployment and inference
- Describe why monitoring is important
- Detect data drifts in the underlying input data
- Demonstrate how to monitor ML models for bias
- Explain how to monitor model resource consumption and latency
- Discuss how to integrate human-in-the-loop reviews of model results in production

Prerequisites

Required

- AWS Technical Essentials course (classroom or digital)
- DevOps Engineering on AWS course, or equivalent experience
- Practical Data Science with Amazon SageMaker course, or equivalent experience

Recommended

- The Elements of Data Science (digital course), or equivalent experience
- Machine Learning Terminology and Process (digital course)

Activities

This course includes presentations, labs, demonstrations, workbooks, and group exercises.

Course Outline

Day 1

Module 0: Welcome

- Course introduction

Module 1: Introduction to MLOps

- Machine learning operations
- Goals of MLOps
- Communication

- From DevOps to MLOps
- ML workflow
- Scope
- MLOps view of ML workflow
- MLOps cases

Module 2: MLOps Development

- Intro to build, train, and evaluate machine learning models
- MLOps security
- Automating
- Apache Airflow
- Kubernetes integration for MLOps
- Amazon SageMaker for MLOps
- Lab: Bring your own algorithm to an MLOps pipeline
- Demonstration: Amazon SageMaker
- Intro to build, train, and evaluate machine learning models
- Lab: Code and serve your ML model with AWS CodeBuild
- Activity: MLOps Action Plan Workbook

Day 2

Module 3: MLOps Deployment

- Introduction to deployment operations
- Model packaging
- Inference
- Lab: Deploy your model to production
- SageMaker production variants
- Deployment strategies
- Deploying to the edge
- Lab: Conduct A/B testing
- Activity: MLOps Action Plan Workbook

Day 3

Module 4: Model Monitoring and Operations

- Lab: Troubleshoot your pipeline
- The importance of monitoring

- Monitoring by design
- Lab: Monitor your ML model
- Human-in-the-loop
- Amazon SageMaker Model Monitor
- Demonstration: Amazon SageMaker Pipelines, Model Monitor, model registry, and Feature Store
- Solving the Problem(s)
- Activity: MLOps Action Plan Workbook

Module 5: Wrap-up

- Course review
- Activity: MLOps Action Plan Workbook
- Wrap-up