

Cloud-Native MLOps with Azure ML and MLflow

Duration: 24 hours

Course Overview:

This course offers a hands-on introduction to MLOps, covering containerization with Docker, experiment tracking with MLflow, CI/CD pipelines, and cloud-native deployment using Azure Machine Learning. Participants will learn to build automated ML workflows, manage models, deploy real-time and batch endpoints, and apply production strategies such as staging, monitoring, and canary releases. Through practical labs, learners will deploy complete ML pipelines from development to production while gaining experience troubleshooting real-world scenarios and applying enterprise best practices.

Course Pre-requisite:

The participant must bring their own Azure Subscription for lab performance.

- Basic understanding of Machine Learning concepts and workflows
- Working knowledge of Python
- Familiarity with Git and version control fundamentals
- Basic awareness of Azure cloud computing concepts
- Understanding of REST APIs and command-line tools.

Course Contents

Module 1: Containerization and Deployment

- Docker fundamentals and advanced concepts
- Containerizing ML workflows and models
- Container registry usage and best practices

Module 2: Experiment Tracking and Model Management

- MLflow fundamentals
- Tracking experiments, metrics, and artifacts
- MLflow model registry and versioning
- Serving models using MLflow

Module 3: CI/CD Pipelines for ML

- CI pipelines for data ingestion and model training
- CD pipelines for model deployment
- GitHub Actions or Azure DevOps pipelines
- Secret management and environment variables

Module 4: Azure ML and Cloud-Native MLOps

- Azure Machine Learning workspace setup
- Managing compute clusters
- Pipelines, jobs, and environments
- Real-time and batch endpoints
- Monitoring, drift detection, and logging

Module 5: Productionizing ML Models

- Staging, testing, and release strategies
- Blue-green and canary deployments
- Feature store introduction
- Model governance, security, and lineage

Hands-On Lab Environment

- Practical labs for each module
- Deploying a complete ML model from development to staging to production
- Building pipelines across Fabric and Azure
- Troubleshooting real-world deployment scenarios