

Training Course on Cloud MLOps on AWS (SageMaker Advanced)

Professional Development Training Course Outline

Category	Data Science	Duration	10 Days
Base Fee	\$3,000 USD	Delivery Mode	Online / On-site Hybrid

Course Introduction

Training Course on Cloud MLOps on AWS (SageMaker Advanced): Deep Dive into AWS Services for MLOps provides a comprehensive, hands-on deep dive into **Cloud MLOps** principles and practices, specifically leveraging advanced **AWS services**, with a strong focus on **Amazon SageMaker**. Participants will master the end-to-end Machine Learning Lifecycle, from data preparation and model experimentation to robust deployment, monitoring, and governance in a production environment. This course emphasizes building scalable, repeatable, and automated MLOps pipelines on the AWS cloud, empowering professionals to accelerate their **AI/ML initiatives** and drive significant business value.

Gain expertise in establishing a resilient MLOps culture and framework, ensuring seamless collaboration between data scientists, ML engineers, and operations teams. Through practical labs and real-world case studies, attendees will learn to implement **CI/CD for ML**, manage **model versioning and lineage**, perform **continuous monitoring for data and model drift**, and optimize **resource utilization** on AWS. This advanced training equips participants with the critical skills to operationalize machine learning at scale, transforming experimental models into reliable, high-performing production systems.

Programme Curriculum

Training Course on Cloud MLOps on AWS (SageMaker Advanced): Deep Dive into AWS Services for MLOps

Introduction

Training Course on Cloud MLOps on AWS (SageMaker Advanced): Deep Dive into AWS Services for MLOps provides a comprehensive, hands-on deep dive into **Cloud MLOps** principles and practices, specifically leveraging advanced **AWS services**, with a strong focus on **Amazon**

SageMaker. Participants will master the end-to-end Machine Learning Lifecycle, from data preparation and model experimentation to robust deployment, monitoring, and governance in a production environment. This course emphasizes building scalable, repeatable, and automated MLOps pipelines on the AWS cloud, empowering professionals to accelerate their **AI/ML initiatives** and drive significant business value.

Gain expertise in establishing a resilient MLOps culture and framework, ensuring seamless collaboration between data scientists, ML engineers, and operations teams. Through practical labs and real-world case studies, attendees will learn to implement **CI/CD for ML**, manage **model versioning and lineage**, perform **continuous monitoring for data and model drift**, and optimize **resource utilization** on AWS. This advanced training equips participants with the critical skills to operationalize machine learning at scale, transforming experimental models into reliable, high-performing production systems.

Course Duration

10 days

Course Objectives

1. Master best practices for building robust, fault-tolerant MLOps pipelines on AWS.
2. Develop proficiency in using AWS services like SageMaker Pipelines and Step Functions for CI/CD in ML.
3. Leverage Amazon S3, AWS Glue, and Amazon SageMaker Feature Store for efficient data versioning, transformation, and feature engineering.
4. Utilize SageMaker Studio and its integrated tools for hyperparameter tuning, distributed training, and experiment tracking.
5. Deploy ML models to production using SageMaker Endpoints, Batch Transform, and SageMaker Inference Recommender for optimal performance.
6. Set up Amazon CloudWatch and SageMaker Model Monitor for real-time detection of data drift, model drift, and performance degradation.
7. Implement SageMaker Model Registry for centralized model version control, lineage tracking, and approval workflows.
8. Build automated pipelines using AWS CodeCommit, CodeBuild, and CodePipeline for seamless ML code integration and deployment.
9. Apply AWS IAM, VPC, and KMS best practices to ensure data and model security and compliance.
10. Diagnose and resolve common MLOps challenges, enhancing system reliability and performance.
11. Explore AWS Lambda, Amazon EKS, and Docker for flexible and cost-effective ML deployments.
12. Understand and apply techniques for model explainability (XAI) and fairness within MLOps on AWS.
13. Identify strategies for managing and reducing infrastructure costs for ML workloads on AWS.

Organizational Benefits

- Drastically reduce the time-to-market for new machine learning models by automating the entire ML lifecycle.
- Ensure consistent, high-performing models in production through robust monitoring, automated retraining, and drift detection.

- Foster seamless communication and collaboration between data scientists, ML engineers, and operations teams with standardized MLOps practices.
- Optimize resource utilization, automate manual tasks, and prevent costly model failures through efficient MLOps implementation on AWS.
- Establish clear audit trails, version control, and security measures for ML models, adhering to regulatory requirements.
- Maximize the value derived from machine learning initiatives by successfully deploying and maintaining models at scale.
- Ensure consistency and reproducibility of ML models, enabling better experimentation and faster iteration cycles.

Target Audience

1. MLOps Engineers.
2. DevOps Engineers.
3. Data Scientists.
4. Machine Learning Engineers
5. Cloud Architects.
6. AI/ML Solution Architects.
7. Data Engineers.
8. Technical Leads & Managers.

Course Outline

Module 1: Introduction to MLOps on AWS

- Understanding the MLOps Landscape: Challenges and Opportunities in Production ML.
- MLOps vs. DevOps: Key Differences and Synergies in the AWS Context.
- The MLOps Maturity Model and its Application to AWS Deployments.
- Overview of Key AWS Services for MLOps (SageMaker, S3, Lambda, Step Functions).
- Case Study: How a retail giant reduced model deployment time by 70% using foundational AWS MLOps practices.

Module 2: Data Management and Feature Engineering for MLOps

- Data Versioning and Governance with Amazon S3 and AWS Lake Formation.
- Building Scalable Data Pipelines with AWS Glue for ETL and Feature Preparation.
- Leveraging Amazon SageMaker Feature Store for Reusable and Shareable Features.
- Data Validation and Quality Checks within the MLOps Pipeline.
- Case Study: A financial institution improved fraud detection model accuracy by 15% through a centralized feature store on SageMaker.

Module 3: Advanced Experimentation and Model Development with SageMaker Studio

- Deep Dive into SageMaker Studio: Environment Setup and Customization.
- Automated Machine Learning (AutoML) with SageMaker Autopilot for Rapid Prototyping.
- Distributed Training Strategies with SageMaker (e.g., Data Parallel, Model Parallel).
- Hyperparameter Tuning and Optimization with SageMaker Hyperparameter Tuning.
- Case Study: An automotive company accelerated model iteration cycles by 40% using SageMaker Studio's integrated experimentation tools.

Module 4: Building Automated ML Pipelines with SageMaker Pipelines

- Designing and Orchestrating End-to-End ML Workflows using SageMaker Pipelines.
- Pipeline Steps: Processing, Training, Model Creation, and Evaluation Components.
- Conditional Logic and Callbacks in SageMaker Pipelines for Dynamic Workflows.
- Managing Pipeline Versions and Rollbacks for Reproducibility.
- Case Study: A healthcare provider automated their medical image analysis pipeline, reducing processing time by 60% with SageMaker Pipelines.

Module 5: Model Registry and Versioning

- Centralized Model Management with Amazon SageMaker Model Registry.
- Registering, Tagging, and Versioning Machine Learning Models.
- Model Lifecycle Management: From Staging to Production Approval.
- Tracking Model Lineage and Metadata for Auditability.
- Case Study: An e-commerce platform improved customer recommendation model governance by 80% through a robust SageMaker Model Registry implementation.

Module 6: CI/CD for Machine Learning Models

- Integrating MLOps Pipelines with AWS CodeCommit, CodeBuild, and CodePipeline.
- Automated Testing Strategies for ML Code, Data, and Models.
- Blue/Green and Canary Deployments for ML Models.
- Implementing Release Gates and Approval Workflows for Production Deployments.
- Case Study: A telecommunications company achieved continuous model delivery, deploying new churn prediction models weekly using AWS Code services.

Module 7: Advanced Model Deployment with SageMaker Endpoints

- Real-time Inference with SageMaker Endpoints: Configuration and Optimization.
- Asynchronous Inference for Large Payloads and Batch Predictions.
- Multi-Model Endpoints and Endpoint Configurations for A/B Testing.
- Using SageMaker Inference Recommender for Optimal Instance Selection.
- Case Study: A media streaming service optimized their content recommendation engine's latency by 25% through fine-tuned SageMaker Endpoints.

Module 8: Batch Inference and Serverless Deployments

- Performing Batch Predictions with SageMaker Batch Transform.
- Serverless ML Inference with AWS Lambda and Amazon API Gateway.
- Integrating SageMaker with AWS Step Functions for Complex Inference Workflows.
- Cost-Effective Batch Processing Strategies on AWS.
- Case Study: A logistics company automated daily route optimization predictions for thousands of deliveries using SageMaker Batch Transform.

Module 9: Model Monitoring and Drift Detection

- Setting up Amazon SageMaker Model Monitor for Data Quality and Model Quality Monitoring.
- Detecting Data Drift and Concept Drift in Production.
- Automated Alerts and Notifications for Model Performance Degradation.
- Root Cause Analysis for Model Failures and Performance Issues.
- Case Study: A financial services firm proactively identified and remediated a 10% drop in fraud detection accuracy due to data drift using SageMaker Model Monitor.

Module 10: Retraining and Continuous Improvement

- Automated Model Retraining Triggers based on Drift or Performance Metrics.
- Implementing Retraining Pipelines with SageMaker and AWS Step Functions.
- Managing Retrained Model Versions and Rollbacks.
- Strategies for Human-in-the-Loop (HITL) for Model Improvement.
- Case Study: An advertising technology company implemented automated retraining for their click-through rate prediction model, improving performance by 5% quarterly.

Module 11: Security and Governance in MLOps

- AWS IAM Best Practices for MLOps Roles and Permissions.
- Network Security: VPC, Security Groups, and Private Endpoints for SageMaker.
- Data Encryption at Rest and in Transit with AWS KMS.
- Compliance Considerations and Audit Trails for ML Workflows.
- Case Study: A government agency achieved strict compliance for their classified document classification model by implementing robust AWS security measures.

Module 12: Cost Optimization for MLOps on AWS

- Strategies for Optimizing SageMaker Instance Types and Pricing Models.
- Leveraging Spot Instances for Cost-Effective Training and Batch Inference.
- Monitoring and Managing AWS Costs for ML Workloads with AWS Cost Explorer.
- Rightsizing and Resource Allocation Best Practices.
- Case Study: A startup reduced their monthly ML infrastructure costs by 30% through intelligent use of SageMaker Spot Instances and reserved instances.

Module 13: Responsible AI and Explainability (XAI)

- Understanding Bias and Fairness in Machine Learning Models.
- Techniques for Model Explainability and Interpretability (e.g., SHAP, LIME).
- Using SageMaker Clarify for Bias Detection and Explainability.
- Building Trustworthy AI Systems in Production.
- Case Study: A hiring platform used SageMaker Clarify to identify and mitigate bias in their resume screening model, leading to more equitable hiring outcomes.

Module 14: Advanced SageMaker Services & Ecosystem Integration

- SageMaker Ground Truth for Scalable Data Labeling.
- Amazon SageMaker Canvas for Business Analysts and Citizen Data Scientists.
- Integrating SageMaker with other AWS Analytics Services (Athena, Redshift).
- Leveraging AWS Marketplace for ML Algorithms and Solutions.
- Case Study: An agricultural tech firm used SageMaker Ground Truth to rapidly label satellite imagery for crop health analysis, accelerating model development.

Module 15: Troubleshooting and Best Practices for Production MLOps

- Common MLOps Challenges and Troubleshooting Techniques on AWS.
- Debugging Failed Pipelines and Model Deployment Issues.
- Monitoring Best Practices for ML Systems: Metrics, Logging, and Dashboards.
- Building Resilient and Self-Healing MLOps Architectures.
- Case Study: A logistics company recovered quickly from a production model failure by implementing advanced logging and automated rollback procedures.

Training Methodology

This course employs a participatory and hands-on approach to ensure practical learning, including:

- Interactive lectures and presentations.
- Group discussions and brainstorming sessions.
- Hands-on exercises using real-world datasets.
- Role-playing and scenario-based simulations.
- Analysis of case studies to bridge theory and practice.
- Peer-to-peer learning and networking.
- Expert-led Q&A sessions.
- Continuous feedback and personalized guidance.

Register as a group from 3 participants for a Discount

Send us an email: info@fineskilltrainingcenter.org or call +254769199797

Certification

Upon successful completion of this training, participants will be issued with a globally- recognized certificate.

Tailor-Made Course

We also offer tailor-made courses based on your needs.

Key Notes

- a. The participant must be conversant with English.
- b. Upon completion of training the participant will be issued with an Authorized Training Certificate
- c. Course duration is flexible and the contents can be modified to fit any number of days.
- d. The course fee includes facilitation training materials, 2 coffee breaks, buffet lunch and A Certificate upon successful completion of Training.
- e. One-year post-training support Consultation and Coaching provided after the course.
- f. Payment should be done at least a week before commence of the training, to Fineskill Training Center account, as indicated in the invoice so as to enable us prepare better for you.

Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
21 Sep 2026	02 Oct 2026	Online	\$2,000
21 Sep 2026	02 Oct 2026	Physical	\$3,000
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0
21 Sep 2026	02 Oct 2026	Physical	\$0

Ready to enroll or request a customized program? Book online or contact us:

Register Online Now

Email: info@fineskilltrainingcenter.com | Website: www.fineskilltrainingcenter.com