

MLA-C01: AWS Machine Learning Engineer – Associate

Take models from notebook to production on AWS with SageMaker — and prepare for the MLA-C01 Machine Learning Engineer exam.

DURATION

3 days · 24 hours

LEVEL

Intermediate

CURRICULUM

9 modules · 6 labs

DELIVERYLive online, Onsite,
Corporate

Course overview

Building a model is only part of the job. Machine learning engineers also prepare the data, choose sensible training approaches, deploy models safely, automate retraining and notice when a model stops performing. This three-day course focuses on that engineering side of machine learning on AWS, using Amazon SageMaker as the main platform.

Day one covers data preparation: ingesting and storing training data, cleaning and transforming it with SageMaker Data Wrangler and Glue, engineering features, handling class imbalance and bias, and managing features in Feature Store. Day two is model development: choosing between built-in algorithms, frameworks and foundation models, training and tuning, evaluating with the right metrics and explaining predictions with SageMaker Clarify.

Day three moves into deployment and operations. You will compare real-time, serverless, asynchronous and batch inference, automate the workflow with SageMaker Pipelines and Model Registry, and set up Model Monitor to catch data and quality drift. Security topics such as IAM roles, VPC isolation and encryption are covered alongside cost controls for training and hosting.

You do not need to be a data scientist to follow the course, but you should be comfortable with Python. The final session reviews exam technique and practice questions. This course is aligned to the MLA-C01 exam and prepares you for it; confirm the current exam guide on the AWS Certification site before booking.

Highlights

- 24 hours of live, hands-on ML engineering
- End-to-end SageMaker workflow from data to endpoint
- MLOps with pipelines, model registry and monitoring
- Mapped to the four MLA-C01 exam domains
- Practice questions with detailed explanations

What's included

- One model taken from raw data to monitored endpoint
- SageMaker Studio, Pipelines and Model Registry in practice
- Bias detection and explainability with Clarify
- Inference options compared on cost and latency
- Drift monitoring with Model Monitor
- Practice questions with explained answers
- Course completion certificate from Ace Skills Consulting

Who this course is for

Who should attend

- ML engineers and data scientists deploying models on AWS
- Software engineers moving into machine learning roles
- MLOps and platform engineers
- Data engineers supporting ML teams

Prerequisites

- Comfortable programming in Python
- Basic understanding of ML concepts such as training, validation and overfitting
- Familiarity with core AWS services such as S3 and IAM

Detailed curriculum — 3 days, 24 hours, 9 modules, 6 hands-on labs

Day 1 Data Preparation for Machine Learning

8 hours

1.1 Data ingestion and storage

- Training data formats and storage options
- Streaming data for ML
- S3, EFS and FSx for Lustre for training data
- Data validation basics

LAB Set up a SageMaker Studio domain and load a dataset · 35 min

1.2 Transformation and feature engineering

- Cleaning, encoding and scaling features
- SageMaker Feature Store
- SageMaker Data Wrangler and Processing jobs
- Handling imbalance and missing values

LAB Engineer features with Data Wrangler and publish to Feature Store · 55 min

1.3 Data quality and bias

- Pre-training bias metrics with Clarify
- Protecting sensitive data
- Labelling with SageMaker Ground Truth

Day 2 Model Development

8 hours

2.1 Choosing an approach

- Built-in algorithms versus script mode
- AWS AI services as alternatives to custom models
- Using foundation models with JumpStart and Amazon Bedrock

2.2 Training and tuning

- Training jobs, instance types and Spot training
- Distributed training at a glance
- Automatic model tuning
- Experiments tracking

LAB Train and tune an XGBoost model with automatic model tuning · 60 min

2.3 Evaluation and explainability

- Classification and regression metrics
- Explainability with SHAP and Clarify
- Detecting overfitting and underfitting
- Model debugging

LAB Evaluate the model and generate an explainability report · 40 min

Day 3 Deployment, MLOps, Monitoring & Exam Readiness

8 hours

3.1 Deployment options

- Real-time, serverless, asynchronous and batch inference
- Auto scaling endpoints
- Multi-model and multi-container endpoints
- Shadow and canary deployments

LAB Deploy a real-time endpoint and a batch transform job · 45 min

3.2 Automation and monitoring

- SageMaker Pipelines
- Model Monitor for data and quality drift
- Model Registry and approval workflows
- CI/CD for ML with CodePipeline

LAB Build a training pipeline with a registry approval step · 55 min

3.3 Security, cost and exam readiness

- IAM roles, VPC isolation and encryption
- Cost controls for training and hosting
- MLA-C01 domain recap and practice set

Skills and tools

Skills you will build

Amazon SageMaker

Feature engineering

SageMaker Feature Store

Hyperparameter tuning

SageMaker Clarify

Model deployment

SageMaker Pipelines

Model Registry

Model Monitor

MLOps

Amazon Bedrock

Tools used in class

Amazon SageMaker Studio

Jupyter notebooks

Python

pandas

scikit-learn

XGBoost

AWS CLI

Certification exam

This course is aligned to the exam below and prepares you for it. The exam itself is booked and taken with the vendor.

Exam	AWS Certified Machine Learning Engineer – Associate
Exam code	MLA-C01
Questions	65
Duration	130 minutes
Passing score	720 / 1000 (scaled)
Delivery	Online proctored, Test centre
Validity	3 years
Official page	https://aws.amazon.com/certification/certified-machine-learning-engineer-associate/

Exam skill areas covered

- Data Preparation for Machine Learning
- ML Model Development
- Deployment and Orchestration of ML Workflows
- ML Solution Monitoring, Maintenance, and Security

Exam details are pending verification. Confirm the current exam code, format and fee on the official vendor page before booking.

Frequently asked questions

Do I need to be a data scientist?

No. You need working Python and a basic grasp of ML concepts; the course focuses on engineering and operations rather than advanced statistics.

Does the course cover generative AI?

It covers using foundation models through SageMaker JumpStart and Amazon Bedrock where they relate to the exam, but the focus is on the ML engineering lifecycle.

Will the labs incur AWS costs?

Labs run in an environment arranged for the batch, and we show you how to shut down resources if you repeat them in your own account.

Does the exam use new question types?

AWS has introduced additional question types such as ordering and matching on some newer exams. Check the current MLA-C01 exam guide for the formats in use.

Can you train our data science team?

Yes. Private batches can focus on the deployment and MLOps patterns your team needs most.

Book a batch or talk to an advisor

Course page: aceskillsconsulting.com/course/mla-c01-aws-machine-learning-engineer-associate · Email: info@aceskillsconsulting.com

Chennai office

Guindy, Chennai, Tamil Nadu

Delhi office

B-95, Old Post Office Street, Shakarpur
New Delhi 110092
Office: 98180 12523

Amazon Web Services, AWS and the related marks are trademarks of Amazon.com, Inc. or its affiliates. All product names, logos, brands, trademarks and registered trademarks are the property of their respective owners. Ace Skills Consulting is an independent training provider; unless explicitly stated, it is not affiliated with, endorsed by or sponsored by any vendor mentioned on this site. Exam details are provided for guidance and may change — always confirm on the vendor's official website. Curriculum generated on 23 Oct 2026; topics and order may be adjusted by the trainer to suit each batch.