Machine Learning Engineering on AWS

Category: AI, Data & Analytics | **Vendor:** Amazon Web Services

Duration: 24.00 hours (3 days) **19.5 CPD Hours** **Rating:** ⭐ 4.6 (5,878 reviews)

Course Information

Delivery Format: Instructor Led - Online

Course Overview

Machine Learning (ML) Engineering on Amazon Web Services (AWS) is a 3-day intermediate course designed for ML professionals seeking to learn machine learning engineering on AWS. Participants learn to build, deploy, orchestrate, and operationalize ML solutions at scale through a balanced combination of theory, practical labs, and activities. Participants will gain practical experience using AWS services such as Amazon SageMaker AI and analytics tools such as Amazon EMR to develop robust, scalable, and production-ready machine learning applications.

About This Course

- Course Introduction
- Introduction to Machine Learning (ML) on AWS
- Analyzing Machine Learning (ML) Challenges
- Data Processing for Machine Learning (ML)
- Data Transformation and Feature Engineering
- Choosing a Modeling Approach
- Training Machine Learning (ML) Models
- Evaluating and Tuning Machine Learning (ML) models
- Model Deployment Strategies
- Securing AWS Machine Learning (ML) Resources
- Machine Learning Operations (MLOps) and Automated Deployment
- Monitoring Model Performance and Data Quality
- Course Wrap-up

Who Should Attend

This course is designed for professionals who are interested in building, deploying, and operationalizing machine learning models on AWS. This could include current and in-training machine learning engineers who might have little prior experience with AWS. Other roles that can benefit from this training are DevOps engineer, developer, and SysOps engineer.

Prerequisites & Entry Requirements

General Prerequisites:

We recommend that attendees of this course have the following:

- Familiarity with basic machine learning concepts
- Working knowledge of Python programming language and common data science libraries such as NumPy, Pandas, and Scikit-learn
- Basic understanding of cloud computing concepts and familiarity with AWS
- Experience with version control systems such as Git (beneficial but not required)

Learning Outcomes

Upon successful completion of this course, participants will be able to:

In this course, you will learn to do the following:

- Explain ML fundamentals and its applications in the AWS Cloud.
- Process, transform, and engineer data for ML tasks by using AWS services.
- Select appropriate ML algorithms and modeling approaches based on problem requirements and model interpretability.
- Design and implement scalable ML pipelines by using AWS services for model training, deployment, and orchestration.
- Create automated continuous integration and delivery (CI/CD) pipelines for ML workflows.
- Discuss appropriate security measures for ML resources on AWS.
- Implement monitoring strategies for deployed ML models, including techniques for detecting data drift.

Detailed Course Outline

Module 0: Course Introduction

Module 1: Introduction to Machine Learning (ML) on AWS

- Topic A: Introduction to ML
- Topic B: Amazon SageMaker AI
- Topic C: Responsible ML

Module 2: Analyzing Machine Learning (ML) Challenges

- Topic A: Evaluating ML business challenges
- Topic B: ML training approaches

- Topic C: ML training algorithms

Module 3: Data Processing for Machine Learning (ML)

- Topic A: Data preparation and types
- Topic B: Exploratory data analysis
- Topic C: AWS storage options and choosing storage

Module 4: Data Transformation and Feature Engineering

- Topic A: Handling incorrect, duplicated, and missing data
- Topic B: Feature engineering concepts
- Topic C: Feature selection techniques
- Topic D: AWS data transformation services
- Lab 1: Analyze and Prepare Data with Amazon SageMaker Data Wrangler and Amazon EMR
- Lab 2: Data Processing Using SageMaker Processing and the SageMaker Python SDK

Module 5: Choosing a Modeling Approach

- Topic A: Amazon SageMaker AI built-in algorithms
- Topic B: Amazon SageMaker Autopilot
- Topic C: Selecting built-in training algorithms
- Topic D: Model selection considerations
- Topic E: ML cost considerations

Module 6: Training Machine Learning (ML) Models

- Topic A: Model training concepts
- Topic B: Training models in Amazon SageMaker AI
- Lab 3: Training a model with Amazon SageMaker AI

Module 7: Evaluating and Tuning Machine Learning (ML) models

- Topic A: Evaluating model performance

- Topic B: Techniques to reduce training time

- Topic C: Hyperparameter tuning techniques

- Lab 4: Model Tuning and Hyperparameter Optimization with Amazon SageMaker AI

Module 8: Model Deployment Strategies

- Topic A: Deployment considerations and target options

- Topic B: Deployment strategies

- Topic C: Choosing a model inference strategy

- Topic D: Container and instance types for inference

- Lab 5: Shifting Traffic

Module 9: Securing AWS Machine Learning (ML) Resources

- Topic A: Access control

- Topic B: Network access controls for ML resources

- Topic C: Security considerations for CI/CD pipelines

Module 10: Machine Learning Operations (MLOps) and Automated Deployment

- Topic A: Introduction to MLOps

- Topic B: Automating testing in CI/CD pipelines

- Topic C: Continuous delivery services

- Lab 6: Using Amazon SageMaker Pipelines and the Amazon SageMaker Model Registry with Amazon SageMaker Studio

Module 11: Monitoring Model Performance and Data Quality

- Topic A: Detecting drift in ML models

- Topic B: SageMaker Model Monitor

- Topic C: Monitoring for data quality and model quality

- Topic D: Automated remediation and troubleshooting

- Lab 7: Monitoring a Model for Data Drift

Module 12: Course Wrap-up, Course Introduction, Introduction to Machine Learning (ML) on AWS, Analyzing Machine Learning (ML) Challenges, Data Processing for Machine Learning (ML), Data Transformation and Feature Engineering, Choosing a Modeling Approach, Training Machine Learning (ML) Models, Evaluating and Tuning Machine Learning (ML) models, Model Deployment Strategies, Securing AWS Machine Learning (ML) Resources, Machine Learning Operations (MLOps) and Automated Deployment, Monitoring Model Performance and Data Quality, Course Wrap-up

Skills You Will Learn:

In this course, you will learn to do the following:

- Explain ML fundamentals and its applications in the AWS Cloud.
- Process, transform, and engineer data for ML tasks by using AWS services.
- Select appropriate ML algorithms and modeling approaches based on problem requirements and model interpretability.
- Design and implement scalable ML pipelines by using AWS services for model training, deployment, and orchestration.
- Create automated continuous integration and delivery (CI/CD) pipelines for ML workflows.
- Discuss appropriate security measures for ML resources on AWS.
- Implement monitoring strategies for deployed ML models, including techniques for detecting data drift.

Frequently Asked Questions

Q: What delivery options are available for Machine Learning Engineering on AWS?

We offer multiple delivery formats:

- Live Instructor-Led Classroom Online (Virtual/Live Online)
- Traditional Instructor-Led Classroom Training (ILT)
- On-site delivery at your offices anywhere in United Kingdom
- Private dedicated courses customized for your team

Q: How many CPD hours does this course provide?

The 3-day Machine Learning Engineering on AWS course provides up to 19.5 CPD hours of structured learning. CPD certificates can be provided upon request.

Q: What is the duration of the Machine Learning Engineering on AWS training?

The training takes place over 3 day(s), with each day lasting approximately 24.00 hours including breaks for lunch and refreshments.

Q: Do you provide corporate training for Machine Learning Engineering on AWS?

Yes, we provide corporate training, dedicated training, and closed classes for Machine Learning Engineering on AWS. Training can take place anywhere in United Kingdom including London, Manchester, Birmingham, Edinburgh, or live online allowing teams from across United Kingdom or internationally to attend.

Q: What related terms do people search for?

Popular related searches include: Machine Learning Engineering on AWS, Machine Learning, MLEA, Amazon Web Services, vILT

Q: Why choose Nexus Human for Machine Learning Engineering on AWS?

Nexus Human is recognized as one of the leading training providers. Our trainers have won multiple awards including:

- Small Firms Best Trainer Award
- National Training Partner of the Year (Ireland) - Multiple Years
- Global Top 30 Instructor Awards (2012, 2019, 2021)
- Tech Excellence Award Nominations
- Learning Performance Institute (LPI) External Training Provider Sponsor 2024

Q: Are there any discount codes available?

Yes! Use discount code **PENPAL5** when booking your Machine Learning Engineering on AWS training. Please note that only one discount code can be used per booking and cannot be combined with other special offers.

Nexus Human

Professional Training & Development

Email: info@nexushuman.com

Website: www.nexushuman.com

Phone: +353 1 889 8700 (Ireland) | +44 20 8142 8961 (UK)

© 2026 Nexus Human. All rights reserved. This brochure was generated on 14/11/2026.