

Machine Learning Engineering on AWS

Duration: 3 Days Course Code: GK910029 Delivery Method: Company Event

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.

Course level: Intermediate

Company Events

These events can be delivered exclusively for your company at our locations or yours, specifically for your delegates and your needs. The Company Events can be tailored or standard course deliveries.

Target Audience:

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.

Objectives:

- 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.
-

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)
-

Content:

Day 1	Lab 2: Data Processing Using SageMaker Processing and the SageMaker Python SDK	Topic 8C: Choosing a model inference strategy
Module 0: Course Introduction	Day 2	Topic 8D: Container and instance types for inference
Module 1: Introduction to Machine Learning (ML) on AWS	Module 5: Choosing a Modeling Approach	Lab 5: Shifting Traffic A/B
Topic 1A: Introduction to ML	Topic 5A: Amazon SageMaker AI built-in algorithms	Day 3
Topic 1B: Amazon SageMaker AI	Topic 5B: Selecting built-in training algorithms	Module 9: Securing AWS Machine Learning (ML) Resources
Topic 1C: Responsible ML	Topic 5C: Amazon SageMaker Autopilot	Topic 9A: Access control
Module 2: Analyzing Machine Learning (ML) Challenges	Topic 5D: Model selection considerations	Topic 9B: Network access controls for ML resources
Topic 2A: Evaluating ML business challenges	Topic 5E: ML cost considerations	Topic 9C: Security considerations for CI/CD pipelines
Topic 2B: ML training approaches	Module 6: Training Machine Learning (ML) Models	Module 10: Machine Learning Operations (MLOps) and Automated Deployment
Topic 2C: ML training algorithms	Topic 6A: Model training concepts	Topic 10A: Introduction to MLOps
Module 3: Data Processing for Machine Learning (ML)	Topic 6B: Training models in Amazon SageMaker AI	Topic 10B: Automating testing in CI/CD pipelines
Topic 3A: Data preparation and types	Lab 3: Training a model with Amazon SageMaker AI	Topic 10C: Continuous delivery services
Topic 3B: Exploratory data analysis	Module 7: Evaluating and Tuning Machine Learning (ML) models	Lab 6: Using Amazon SageMaker Pipelines and the Amazon SageMaker Model Registry with Amazon SageMaker Studio
Topic 3C: AWS storage options and choosing storage	Topic 7A: Evaluating model performance	Module 11: Monitoring Model Performance and Data Quality
Module 4: Data Transformation and Feature Engineering	Topic 7B: Techniques to reduce training time	Topic 11A: Detecting drift in ML models
Topic 4A: Handling incorrect, duplicated, and missing data	Topic 7C: Hyperparameter tuning techniques	Topic 11B: SageMaker Model Monitor
Topic 4B: Feature engineering concepts	Lab 4: Model Tuning and Hyperparameter Optimization with Amazon SageMaker AI	Topic 11C: Monitoring for data quality and model quality
Topic 4C: Feature selection techniques	Module 8: Model Deployment Strategies	Topic 11D: Automated remediation and
Topic 4D: AWS data transformation services	Topic 8A: Deployment considerations and target options	

Lab 1: Analyze and Prepare Data with Amazon SageMaker Data Wrangler and Amazon EMR

Topic 8B: Deployment strategies

troubleshooting

Lab 7: Monitoring a Model for Data Drift

Module 12: Course Wrap-up

Additional Information:

Activities

This course includes presentations, hands-on labs, demonstrations, and group exercises.

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/