

Description

This full-day, intermediate-level course provides a technical introduction to machine learning (ML) on AWS. It emphasizes the ML expertise required for consulting partners to develop ML practices and independent software vendors (ISVs) to support ML services. This course introduces you to the full Amazon Machine Learning (Amazon ML) stack and describes the service, capabilities, and limitations. You will learn ML design principles, requirements, and data dependencies. You will also apply those principles through the deployment of a custom machine learning model using Amazon SageMaker. Finally, the course identifies key use cases for ML technology and prepares you to engage customers in ML solutions discussions.

Intended Audience

This course is intended for:

- Technical professionals at APN Consulting and Technology Partner organizations involved in machine learning (ML)
- Developers, data scientists, and technical account managers focused on ML

Course Objectives

In this course, you will learn how to:

- Engage customers in machine learning discussions and qualify opportunities
- Identify key use cases for machine learning
- Explain the design process, requirements, and data dependencies for implementing ML on AWS
- Describe the capabilities of ML managed services on AWS
- Create and deploy a custom machine learning model at scale using Amazon SageMaker
- Describe the end-to-end process of implementing machine learning features

Prerequisites

We recommend that attendees of this course have the following:

- *AWS Technical Professional*
- Familiarity with ML concepts
- Familiarity with cloud computing concepts
- Experience with Python

Delivery Method

This course is delivered through:

- Classroom training

Note: There are no technical lab exercises in this course.

Duration

Full day (eight hours)

Course Outline

This course covers the following concepts:

- Introduction to machine learning
- Machine learning processes
- Deep learning with Amazon Machine Images (AMIs)
- ML managed services on AWS

- Amazon SageMaker concepts
- Amazon QuickSight ML
- Cost considerations and optimization
- Continuous integration and deployment
- Machine learning solutions
- How to build a machine learning practice
- How to incorporate ML into decision making processes
- Amazon SageMaker Marketplace
- Machine learning with Amazon Elastic Compute Cloud (EC2) instance types
- TensorFlow optimization
- Amazon SageMaker topics:
 - Ground truth
 - Software-defined data center notebooks
 - Built-in algorithms
 - Graphical user interface (GUI) for training and deployment
 - Local
 - Inference
 - Hyperparameter optimization
 - Batch transform
 - Reinforcement learning
 - Edge