

Introduction

The AWS Certified Machine Learning—Specialty (MLS-C01) exam is intended for individuals who perform a Development or Data Science role. This exam validates an examinee’s ability to build, train, tune, and deploy machine learning (ML) models using the AWS Cloud.

It validates an examinee’s ability to design, implement, deploy, and maintain ML solutions for given business problems. It will validate the candidate’s ability to:

- Select and justify the appropriate ML approach for a given business problem.
- Identify appropriate AWS services to implement ML solutions.
- Design and implement scalable, cost-optimized, reliable, and secure ML solutions.

Recommended AWS Knowledge

The successful candidate likely has one to two years of hands-on experience developing, architecting, or running ML/deep learning workloads on the AWS Cloud, along with:

- The ability to express the intuition behind basic ML algorithms
- Experience performing basic hyperparameter optimization
- Experience with ML and deep learning frameworks
- The ability to follow model-training best practices
- The ability to follow deployment and operational best practices

Exam Preparation

These training courses and materials may be helpful for examination preparation:

AWS Training (aws.amazon.com/training)

- Big Data on AWS: 3-day instructor-led live [course](#)
- Deep Learning on AWS: 1-day instructor-led live [course](#)

Other Materials

- Machine Learning on AWS [web-page](#)
- Amazon Machine Learning Concepts [document](#)
- Splitting Your Data [document](#)
- Types of Machine Learning Models [document](#)
- Data Transformations [document](#)
- Containerized Machine Learning on AWS [video](#)
- AWS re:Invent Machine Learning talk [video](#)

Exam Content

Response Types

There are two types of questions on the examination:

- **Multiple-choice:** Has one correct response and three incorrect responses (distractors).
- **Multiple-response:** Has two or more correct responses out of five or more options.

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that an examinee with incomplete knowledge or skill would likely choose. However, they are generally plausible responses that fit in the content area defined by the test objective.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

Unscored Content

Your examination may include unscored items that are placed on the test to gather statistical information. These items are not identified on the form and do not affect your score.

Exam Results

The AWS Certified Machine Learning—Specialty (MLS-C01) exam is a pass or fail exam. The examination is scored against a minimum standard established by AWS professionals who are guided by certification industry best practices and guidelines.

Your results for the examination are reported as a score from 100-1000, with a minimum passing score of 750. Your score shows how you performed on the examination as a whole and whether or not you passed. Scaled scoring models are used to equate scores across multiple exam forms that may have slightly different difficulty levels.

Your score report contains a table of classifications of your performance at each section level. This information is designed to provide general feedback concerning your examination performance. The examination uses a compensatory scoring model, which means that you do not need to “pass” the individual sections, only the overall examination. Each section of the examination has a specific weighting, so some sections have more questions than others. The table contains general information, highlighting your strengths and weaknesses. Exercise caution when interpreting section-level feedback.

Content Outline

This exam guide includes weightings, test domains, and objectives only. It is not a comprehensive listing of the content on this examination. The table below lists the main content domains and their weightings.

| Domain | % of Examination |
|--|------------------|
| Domain 1: Data Engineering | 20% |
| Domain 2: Exploratory Data Analysis | 24% |
| Domain 3: Modeling | 36% |
| Domain 4: Machine Learning Implementation and Operations | 20% |
| TOTAL | 100% |

Domain 1: Data Engineering

- 1.1 Create data repositories for machine learning.
- 1.2 Identify and implement a data-ingestion solution.
- 1.3 Identify and implement a data-transformation solution.

Domain 2: Exploratory Data Analysis

- 2.1 Sanitize and prepare data for modeling.
- 2.2 Perform feature engineering.
- 2.3 Analyze and visualize data for machine learning.

Domain 3: Modeling

- 3.1 Frame business problems as machine learning problems.
- 3.2 Select the appropriate model(s) for a given machine learning problem.
- 3.3 Train machine learning models.
- 3.4 Perform hyperparameter optimization.
- 3.5 Evaluate machine learning models.

Domain 4: Machine Learning Implementation and Operations

- 4.1 Build machine learning solutions for performance, availability, scalability, resiliency, and fault tolerance.
- 4.2 Recommend and implement the appropriate machine learning services and features for a given problem.
- 4.3 Apply basic AWS security practices to machine learning solutions.
- 4.4 Deploy and operationalize machine learning solutions.