

## AWS Certified Machine Learning – Specialty (MLS-C01) Exam Guide

### Introduction

The AWS Certified Machine Learning – Specialty (MLS-C01) exam is intended for individuals who perform an artificial intelligence/machine learning (AI/ML) development or data science role. The exam validates a candidate's ability to design, build, deploy, optimize, train, tune, and maintain ML solutions for given business problems by using the AWS Cloud.

The exam also validates a candidate's ability to complete the following tasks:

- 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

### Target candidate description

The target candidate is expected to have 2 or more years of hands-on experience developing, architecting, and running ML or deep learning workloads in the AWS Cloud.

#### Recommended AWS knowledge

The target candidate should have the following knowledge:

- 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 best practices
- The ability to follow operational best practices

#### What is considered out of scope for the target candidate?

The following is a non-exhaustive list of related job tasks that the target candidate is not expected to be able to perform. These items are considered out of scope for the exam:

- Extensive or complex algorithm development
- Extensive hyperparameter optimization
- Complex mathematical proofs and computations
- Advanced networking and network design
- Advanced database, security, and DevOps concepts
- DevOps-related tasks for Amazon EMR

For a detailed list of specific tools and technologies that might be covered on the exam, as well as lists of in-scope and out-of-scope AWS services, refer to the Appendix.

## Exam content

### Response types

There are two types of questions on the exam:

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

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing. The exam includes 50 questions that will affect your score.

### Unscored content

The exam includes 15 unscored questions that do not affect your score. AWS collects information about candidate performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

### Exam results

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

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether or not you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Your score report could contain a table of classifications of your performance at each section level. This information is intended to provide general feedback about your exam performance. The exam uses a compensatory scoring model, which means that you do not need to achieve a passing score in each section. You need to pass only the overall exam.

Each section of the exam has a specific weighting, so some sections have more questions than other sections have. The table contains general information that highlights your strengths and weaknesses. Use caution when interpreting section-level feedback.

### Content outline

This exam guide includes weightings, test domains, and objectives for the exam. It is not a comprehensive listing of the content on the exam. However, additional context for each of the objectives is available to help guide your preparation for the exam. The following table lists the main content domains and their weightings. The table precedes the complete exam content outline, which includes the additional context. The percentage in each domain represents only scored content.

Domain	% of Exam
Domain 1: Data Engineering	20%
Domain 2: Exploratory Data Analysis	24%
Domain 3: Modeling	36%
Domain 4: Machine Learning Implementation and Operations	20%
<b>TOTAL</b>	<b>100%</b>

## Domain 1: Data Engineering

- 1.1 Create data repositories for machine learning.
  - Identify data sources (e.g., content and location, primary sources such as user data)
  - Determine storage mediums (e.g., DB, Data Lake, S3, EFS, EBS)
- 1.2 Identify and implement a data ingestion solution.
  - Data job styles/types (batch load, streaming)
  - Data ingestion pipelines (Batch-based ML workloads and streaming-based ML workloads)
    - Kinesis
    - Kinesis Analytics
    - Kinesis Firehose
    - EMR
    - Glue
  - Job scheduling
- 1.3 Identify and implement a data transformation solution.
  - Transforming data transit (ETL: Glue, EMR, AWS Batch)
  - Handle ML-specific data using map reduce (Hadoop, Spark, Hive)

## Domain 2: Exploratory Data Analysis

- 2.1 Sanitize and prepare data for modeling.
  - Identify and handle missing data, corrupt data, stop words, etc.
  - Formatting, normalizing, augmenting, and scaling data
  - Labeled data (recognizing when you have enough labeled data and identifying mitigation strategies [Data labeling tools (Mechanical Turk, manual labor)])
- 2.2 Perform feature engineering.
  - Identify and extract features from data sets, including from data sources such as text, speech, image, public datasets, etc.
  - Analyze/evaluate feature engineering concepts (binning, tokenization, outliers, synthetic features, 1 hot encoding, reducing dimensionality of data)
- 2.3 Analyze and visualize data for machine learning.
  - Graphing (scatter plot, time series, histogram, box plot)
  - Interpreting descriptive statistics (correlation, summary statistics, p value)
  - Clustering (hierarchical, diagnosing, elbow plot, cluster size)

## Domain 3: Modeling

- 3.1 Frame business problems as machine learning problems.
  - Determine when to use/when not to use ML
  - Know the difference between supervised and unsupervised learning
  - Selecting from among classification, regression, forecasting, clustering, recommendation, etc.
- 3.2 Select the appropriate model(s) for a given machine learning problem.
  - Xgboost, logistic regression, K-means, linear regression, decision trees, random forests, RNN, CNN, Ensemble, Transfer learning
  - Express intuition behind models
- 3.3 Train machine learning models.
  - Train validation test split, cross-validation
  - Optimizer, gradient descent, loss functions, local minima, convergence, batches, probability, etc.
  - Compute choice (GPU vs. CPU, distributed vs. non-distributed, platform [Spark vs. non-Spark])
  - Model updates and retraining
    - Batch vs. real-time/online
- 3.4 Perform hyperparameter optimization.
  - Regularization
    - Drop out
    - L1/L2
  - Cross validation
  - Model initialization
  - Neural network architecture (layers/nodes), learning rate, activation functions
  - Tree-based models (# of trees, # of levels)
  - Linear models (learning rate)
- 3.5 Evaluate machine learning models.
  - Avoid overfitting/underfitting (detect and handle bias and variance)
  - Metrics (AUC-ROC, accuracy, precision, recall, RMSE, F1 score)
  - Confusion matrix
  - Offline and online model evaluation, A/B testing
  - Compare models using metrics (time to train a model, quality of model, engineering costs)
  - Cross validation

## Domain 4: Machine Learning Implementation and Operations

4.1 Build machine learning solutions for performance, availability, scalability, resiliency, and fault tolerance.

- AWS environment logging and monitoring
  - CloudTrail and CloudWatch
  - Build error monitoring
- Multiple regions, Multiple AZs
- AMI/golden image
- Docker containers
- Auto Scaling groups
- Rightsizing
  - Instances
  - Provisioned IOPS
  - Volumes
- Load balancing
- AWS best practices

4.2 Recommend and implement the appropriate machine learning services and features for a given problem.

- ML on AWS (application services)
  - Poly
  - Lex
  - Transcribe
- AWS service limits
- Build your own model vs. SageMaker built-in algorithms
- Infrastructure: (spot, instance types), cost considerations
  - Using spot instances to train deep learning models using AWS Batch

4.3 Apply basic AWS security practices to machine learning solutions.

- IAM
- S3 bucket policies
- Security groups
- VPC
- Encryption/anonymization

4.4 Deploy and operationalize machine learning solutions.

- Exposing endpoints and interacting with them
- ML model versioning
- A/B testing
- Retrain pipelines
- ML debugging/troubleshooting
  - Detect and mitigate drop in performance
  - Monitor performance of the model

## Appendix

### Which key tools, technologies, and concepts might be covered on the exam?

The following is a non-exhaustive list of the tools and technologies that could appear on the exam. This list is subject to change and is provided to help you understand the general scope of services, features, or technologies on the exam. The general tools and technologies in this list appear in no particular order. AWS services are grouped according to their primary functions. While some of these technologies will likely be covered more than others on the exam, the order and placement of them in this list is no indication of relative weight or importance:

- Ingestion/Collection
- Processing/ETL
- Data analysis/visualization
- Model training
- Model deployment/inference
- Operational
- AWS ML application services
- Language relevant to ML (for example, Python, Java, Scala, R, SQL)
- Notebooks and integrated development environments (IDEs)

### AWS services and features

Analytics:

- Amazon Athena
- Amazon EMR
- Amazon Kinesis Data Analytics
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon QuickSight

Compute:

- AWS Batch
- Amazon EC2

Containers:

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Database:

- AWS Glue
- Amazon Redshift

Internet of Things (IoT):

- AWS IoT Greengrass

#### Machine Learning:

- Amazon Comprehend
- AWS Deep Learning AMIs (DLAMI)
- AWS DeepLens
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Lex
- Amazon Polly
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

#### Management and Governance:

- AWS CloudTrail
- Amazon CloudWatch

#### Networking and Content Delivery:

- Amazon VPC

#### Security, Identity, and Compliance:

- AWS Identity and Access Management (IAM)

#### Serverless:

- AWS Fargate
- AWS Lambda

#### Storage:

- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3

## **Out-of-scope AWS services and features**

The following is a non-exhaustive list of AWS services and features that are not covered on the exam. These services and features do not represent every AWS offering that is excluded from the exam content. Services or features that are entirely unrelated to the target job roles for the exam are excluded from this list because they are assumed to be irrelevant.

Out-of-scope AWS services and features include the following:

- AWS Data Pipeline
- AWS DeepRacer
- Amazon Machine Learning (Amazon ML)