AWS Certified Machine Learning - Specialty (MLS-C01) Prep

AWS MLS-C01 Machine Learning Specialty. SageMaker, data engineering, and model deployment.

IT Certifications Instructor: Brian Taylor • 5 sections • 45 lessons • 15.25h total

View this course online Browse all courses

About this course

This course prepares you for the AWS Certified Machine Learning Specialty (MLS-C01) exam by focusing on the practical application of ML concepts on AWS. We cover SageMaker, data engineering, model training, and deployment.

You will work through hands-on labs, real case studies, and practice exams that mirror the certification test. The course includes video lessons and downloadable study guides.

Course curriculum

Foundations: The AWS ML Ecosystem & Exam Blueprint

- Decoding the MLS-C01 Exam Domains (14m)
  We'll break down the four exam domains, their weightings, and the core skills you need to demonstrate. Includes a downloadable study plan template.
- Core AWS Services for Machine Learning (18m)
  A guided tour of the essential AWS ML services like SageMaker, Rekognition, Comprehend, and how they fit into different ML paradigms.
- Setting Up Your AWS Sandbox (8m)
  Step-by-step guide to configuring a safe and cost-effective AWS environment for your hands-on labs and practice.
- Understanding the ML Lifecycle on AWS (22m)
  Learn the end-to-end machine learning project lifecycle, from data collection to model monitoring, within the AWS context.
- Key IAM and Security Concepts for ML (12m)
  A practical overview of IAM roles, policies, and security best practices specific to managing ML workloads.
- Choosing the Right AWS ML Service for the Job (25m)
  Problem-solving session: We'll analyze business problems and map them to the correct AWS service (e.g., built-in vs. custom models).
- Cost Management for ML Workloads (11m)
  Learn how to estimate costs and optimize spending for data storage, compute, and model hosting on AWS.

Data Engineering: Feeding Your Models

- Architecting Data Ingestion Pipelines with AWS (28m)
  Real case study analysis: Building scalable data pipelines using Kinesis, Glue, and S3 for real-time and batch processing.
- Data Storage Strategies for ML: S3, Athena, and Redshift (19m)
  We'll compare different data storage options and discuss best practices for organizing data for ML training.
- Data Cleaning and Transformation with AWS Glue (24m)
  Step-by-step guide to using AWS Glue for ETL, including handling missing values and feature engineering basics.
- Feature Engineering and Management on SageMaker (26m)
  Learn how to use SageMaker Processing Jobs and Feature Store to create, track, and serve high-quality features.
- Handling Large Datasets and Distributed Data Processing (32m)
  Deep dive into processing massive datasets using Spark on EMR or within SageMaker, a common exam scenario.

- Securing Your Data Pipeline: Encryption and Access Control (15m)
  A focused lesson on implementing encryption at rest and in transit, and managing access for data pipelines.
- Automating Data Validation with SageMaker Data Wrangler (12m)
  Practical walkthrough of using SageMaker Data Wrangler to automate data quality checks and analysis.
- Problem-Solving: Optimizing Data for Model Training (21m)
  We'll analyze a scenario where data format and location impact training performance and choose the optimal solution.
- Real-World Data Challenge: From Raw to Ready (30m)
  A hands-on lab where you'll take a messy dataset and build a complete data prep pipeline using AWS services.

Model Building: Training and Optimization

- Training Models with SageMaker Built-in Algorithms (23m)
  Step-by-step guide to using algorithms like XGBoost and Linear Learner for common ML problems like classification.
- Bringing Your Own Script: Custom Containers with SageMaker (29m)
  Learn how to package your custom training code (e.g., PyTorch, TensorFlow) into a Docker container for SageMaker.
- Hyperparameter Tuning with SageMaker Automatic Model Tuning (17m)
  We'll explore how to set up and manage hyperparameter tuning jobs to find the best model configuration efficiently.
- Training Models at Scale with Distributed Training (27m)
  Understand when and how to use SageMaker's distributed training features to speed up training on large datasets.
- Training, Validation, and Test Sets: Best Practices (13m)
  A conceptual but practical lesson on splitting your data correctly to avoid overfitting and ensure model generalizability.
- Debugging and Profiling Training Jobs with SageMaker Debugger (20m)
  Learn to use SageMaker Debugger to identify issues like vanishing gradients and bottlenecks during training.
- Choosing the Right Instance Type for Training (16m)
  Problem-solving session: We'll match model complexity and data size to the most cost-effective EC2 instance types.
- Transfer Learning and Using Pre-trained Models (18m)
  Explore how to leverage built-in algorithms and pre-trained models from AWS Marketplace to accelerate development.

Model Deployment and MLOps: Going Live

- Deploying Models to SageMaker Endpoints (22m)
  A complete walkthrough of deploying a trained model to a real-time endpoint for inference.

- A/B Testing and Shadow Deployments in SageMaker (25m)
  Learn how to safely roll out new models using SageMaker's deployment strategies to minimize risk.
- Building Batch Transform Pipelines for Inference (19m)
  Step-by-step guide to setting up batch transform jobs for large-scale, offline predictions.
- Optimizing Inference: Cost and Performance (28m)
  We'll cover techniques like model compilation with SageMaker Neo and choosing the right instance types for inference.
- Automating ML Workflows with SageMaker Pipelines (31m)
  Real case study analysis: Building a repeatable, automated pipeline for data prep, training, and deployment.
- Monitoring Model Performance and Drift (23m)
  Learn to use SageMaker Model Monitor to detect data and concept drift in production.
- CI/CD for Machine Learning Models (26m)
  Problem-solving session: Architecting a CI/CD pipeline for your ML models using CodePipeline and CodeBuild.
- Using SageMaker Inference Recommender (10m)
  A quick guide to using the Inference Recommender to find the best deployment configuration for your model.
- Edge Deployment: SageMaker Edge Manager (15m)
  Introduction to deploying and managing models on edge devices, a key topic for IoT scenarios.
- Securing Your Endpoint and Data (14m)
  Deep dive into VPC configurations, endpoint policies, and data encryption for deployed models.
- Hands-On Lab: End-to-End Deployment (35m)
  A comprehensive lab where you'll train a model, deploy it, test it, and clean up all resources.

Advanced Topics and Exam Strategy

- Deep Dive: Computer Vision with Amazon Rekognition (16m)
  We'll explore Rekognition's custom labels feature and when to use it versus building a custom CV model.
- Deep Dive: NLP with Amazon Comprehend (18m)
  Understanding Comprehend's capabilities for sentiment analysis, entity recognition, and custom classification.
- Building Recommendation Engines on AWS (21m)
  A practical look at building personalization engines using Personalize and other SageMaker solutions.
- Time Series Forecasting with SageMaker (24m)
  Learn to use SageMaker's DeepAR and other forecasting algorithms for demand planning and forecasting.
- Machine Learning for Anomaly Detection (19m)
  Problem-solving session: Identifying fraudulent activity or system failures using SageMaker and other AWS tools.
- Using AWS AI Services: A Practical Overview (17m)
  When to use turnkey AI services (Lex, Polly, Translate) versus building custom ML solutions.

- Ethical AI and Fairness in AWS ML (12m)
  A critical discussion on bias detection, explainability (SageMaker Clarify), and responsible AI practices.
- Final Exam Simulation and Question Walkthrough (38m)
  A timed mock exam with detailed explanations for each question to solidify your knowledge and identify weak spots.
- Common Exam Pitfalls and How to Avoid Them (15m)
  We'll review tricky topics, common misinterpretations of exam questions, and key service limits to remember.
- Your 24-Hour Pre-Exam Plan (9m)
  A final checklist and actionable plan for the day before and the day of your exam to maximize your confidence.

Discover our blog

Learn more with free articles written by our professors and guest experts.

Latest article: **Flight Attendant Preparation: What to Study Before You Apply**

Visit the blog All courses This course