

AWS Certified Generative AI Developer (AIP-C01) Prep

AWS AIP-C01 Generative AI Developer. Bedrock, SageMaker, and prompt engineering.

IT Certifications Instructor: Brian Taylor • 5 sections • 47 lessons • 14.50h total

[View this course online](#) [Browse all courses](#)

About this course

This course prepares you for the AWS Certified Generative AI Developer (AIP-C01) exam. We cover Bedrock, SageMaker, prompt engineering, and deploying generative AI solutions on AWS. Hands-on labs, real scenarios, and practice exams included.

Course curriculum

Foundations and Exam Blueprint

- **Decoding the AIP-C01 Exam Structure (12m)**
We'll break down the official exam guide, domains, and question formats to set your study strategy from day one.
- **Core AWS AI Services Overview (18m)**
Learn the key players: Amazon Bedrock, SageMaker, and their supporting services. Understand what each one does and when to use it.
- **Setting Up Your AWS Exam Environment (8m)**
A practical, step-by-step guide to configuring your AWS account for the labs, including IAM roles and budget controls.
- **Essential Machine Learning Concepts for Devs (25m)**
A quick refresher on ML fundamentals, model types, and evaluation metrics you'll need to know for the exam.
- **Understanding Generative AI Architectures (15m)**
We'll cover the high-level architecture for GenAI apps, including RAG, fine-tuning, and agent-based systems.
- **Ethical AI and Responsible Design on AWS (11m)**
Learn AWS's approach to fairness, transparency, and privacy in AI, a key topic for the 'Responsible AI' domain.
- **Downloadable: AIP-C01 Study Planner (7m)**
Get our customizable spreadsheet to track your progress across all exam domains and schedule your practice tests.

Core Concepts: Model Selection and Deployment

- **Choosing Foundational Models in Amazon Bedrock (22m)**
A practical guide to evaluating and selecting the right FM for your use case, including model parameters and providers.
- **Step-by-Step: Your First API Call with Bedrock (14m)**
Hands-on session to invoke a model, parse the response, and handle errors using the AWS SDK.
- **Deploying Custom Models with SageMaker (30m)**
We'll walk through deploying a model from Hugging Face or your own training job to a SageMaker endpoint.
- **Real Case Study: Building a Content Generator (18m)**
Analyze a real-world scenario for a marketing content generator, covering prompt engineering and basic guardrails.
- **Prompt Engineering Techniques for Bedrock (24m)**
Learn zero-shot, few-shot, and chain-of-thought prompting with downloadable templates you can use immediately.

- **Problem-Solving: Handling Model Hallucinations (12m)**
A focused session on techniques to reduce hallucinations, including citation requirements and source grounding.
- **Configuring Inference Parameters for Optimal Results (16m)**
Deep dive into temperature, top_p, and max_tokens to control model output for different applications.
- **Managing Model Versions and Arn Identifiers (9m)**
A practical look at how AWS versions its models and how to correctly reference them in your code and IaC.
- **Lab: Building a Multi-Modal Chatbot (35m)**
Apply what you've learned by building a chatbot that can understand both text and images using Bedrock.
- **Cost Optimization for GenAI Workloads (13m)**
Learn how to manage and forecast costs for model invocation, training, and hosting on AWS.

Deep Dive: RAG, Fine-Tuning, and Agents

- **Architecting a RAG System on AWS (28m)**
We'll design a full RAG pipeline, from document ingestion with OpenSearch to retrieval and generation.
- **Hands-On: Data Ingestion and Chunking Strategies (21m)**
A practical session on processing your documents, choosing chunk sizes, and storing embeddings effectively.
- **Implementing Agents for Complex Task Solving (26m)**
Learn how to use Amazon Bedrock Agents to orchestrate multi-step tasks and integrate with external APIs.
- **Real Case Study: Enterprise Knowledge Base (19m)**
Analyze the architecture for a secure, internal knowledge base using RAG and IAM-based access controls.
- **When to Fine-Tune vs. RAG: A Decision Framework (15m)**
A problem-solving session to help you choose the right approach for your specific business problem.
- **Step-by-Step: Fine-Tuning a Titan Model (32m)**
A detailed walkthrough of preparing a training dataset and starting a fine-tuning job in SageMaker.
- **Evaluating RAG System Performance (17m)**
Learn the key metrics for RAG, including retrieval accuracy and faithfulness, with downloadable evaluation sheets.

Advanced Application: Security, MLOps, and Scaling

- **Securing Your GenAI Application with IAM (23m)**
A deep dive into IAM policies for Bedrock and SageMaker, ensuring least-privilege access for your applications.
- **Protecting Data: PII Redaction and Prompt Filtering (19m)**
Hands-on guide to using AWS native tools to detect and filter sensitive information in prompts and responses.

- **Building a CI/CD Pipeline for GenAI Models (27m)**
We'll use CodePipeline and SageMaker to automate the testing and deployment of a new model version.
- **Monitoring Model Drift and Performance (20m)**
Set up CloudWatch alarms and custom metrics to detect when your model's performance starts to degrade.
- **Problem-Solving: Handling API Throttling and Retries (10m)**
A focused session on building resilient applications that can handle Bedrock API limits gracefully.
- **Scaling Inference with SageMaker Endpoints (24m)**
Learn how to configure auto-scaling for your model endpoints to handle variable traffic loads.
- **Real Case Study: Deploying a Private Model (16m)**
Walk through the architecture for a model that runs in your VPC with no public internet access.
- **Using Guardrails for Safe Application Logic (22m)**
Implement Amazon Bedrock Guardrails to filter topics and deny unwanted content in your app's responses.
- **Downloadable: Security Checklist for GenAI Apps (8m)**
A comprehensive checklist covering IAM, VPC, data encryption, and prompt security for your projects.
- **Orchestration with AWS Step Functions and GenAI (29m)**
Learn to build complex, multi-step GenAI workflows using Step Functions to coordinate different services.
- **Cost and Performance Trade-offs in Production (15m)**
A practical discussion on balancing latency, cost, and accuracy for production-grade GenAI systems.

Implementation, Practice, and Next Steps

- **AIP-C01 Practice Exam: Domain 1 (Model Selection) (25m)**
Test your knowledge with a timed quiz focused on model selection, deployment, and pricing. Includes detailed answer explanations.
- **AIP-C01 Practice Exam: Domain 2 (RAG & Agents) (28m)**
A second timed quiz covering RAG architecture, data ingestion, and agent implementation. We'll review the answers together.
- **AIP-C01 Practice Exam: Domain 3 (Security & MLOps) (26m)**
The final practice exam, focusing on security best practices, monitoring, and deployment strategies.
- **Full Mock Exam Walkthrough (Timed Simulation) (40m)**
A full-length mock exam to simulate the real test environment. We'll review strategies for pacing and tricky questions.
- **Common Exam Pitfalls and How to Avoid Them (13m)**
A review of the most commonly missed question types and concepts on the AIP-C01 exam.
- **Final Review: Key Services and CLI Commands (18m)**
A rapid-fire review of the most important AWS CLI commands and service limits you need to memorize.
- **Your Post-Certification Career Path (12m)**
What to do after you pass: building a portfolio, contributing to projects, and specializing in GenAI.

- **Setting Up Your First Production GenAI Project (21m)**
A step-by-step guide to moving from a prototype to a production-ready application, focusing on the AWS Well-Architected Framework.
- **Staying Current: Resources for 2026 and Beyond (9m)**
A curated list of blogs, newsletters, and AWS re:Invent sessions to keep your skills sharp after the exam.
- **Congratulations and What's Next? (5m)**
A final message from the instructor, plus a guide on how to claim your CPE credits and share your achievement.
- **Downloadable: AIP-C01 Exam Day Checklist (6m)**
A last-minute checklist to ensure you're prepared for exam day, from technical setup to mental readiness.
- **Course Wrap-Up and Key Takeaways (8m)**
A summary of the entire course, highlighting the core concepts and practical skills you've gained.

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](#)