Generative AI Essentials on AWS

AWS Classroom Training

Generative AI Essentials on AWS v3.0

SKU: ILT-TF-100-MLGAIE-30-EN

Course description

Generative AI Essentials on AWS v3.0 is a one-day, instructor-led course for practitioners who already use AI in their daily work and understand the basic concepts. It moves past "what is generative AI?" to the next question practitioners face: how to use it responsibly, build with it, and govern it at scale. The day is organized as two connected arcs. In the Trust Arc, learners work inside AWS Quick as a consumer of AI output and learn to evaluate, understand, fix, and stress-test that output so they can judge its trustworthiness rather than accept it at face value. In the Adoption Arc, learners become the producer: they build a multi-tool agent in Amazon Bedrock AgentCore and widen its audience from personal use to a governed, external-facing service. Responsible AI is a foundational theme throughout, connecting ethical principles to technical practice. Roughly half of the course is hands-on, delivered through two cumulative labs that each build capability across four progressive mindsets. Learners leave able to make informed architecture, governance, and adoption decisions for their teams.

- **Course level:** Fundamental
- **Duration:** 1 day (approximately 7 hours of core content plus breaks)

Activities

This course includes presentations, demonstrations, hands-on labs, and group discussion. Approximately half of course time is dedicated to two cumulative hands-on labs.

Course objectives

In this course, you will learn to do the following:

- Explain what a foundation model can and cannot do on its own, and how a companion tool like AWS Quick closes those gaps.
- Apply prompt engineering techniques in AWS Quick, including grounding, bias mitigation, and data-boundary discipline, to reduce the three output-failure modes and produce trustworthy, defensible output.
- Identify the data sources and access boundaries that shape retrieval, and explain how RAG grounds an answer, distinguishing retrieval from generation and interpreting citations.
- Distinguish companion AI from agentic AI through planning, reasoning, tool use, and multi-step action, and interpret an orchestration trace as the record of what an agent did.
- Decide when and why to move an agent beyond personal use, and deploy it to a wider audience using Amazon Bedrock AgentCore.

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

- Explain layered governance for a hosted agent, including system-prompt scoping, tool restrictions, and guardrails, distinguishing instructions from enforced controls.

Intended audience

This course is intended for the following practitioners who already use AI in their work:

- Business analysts
- IT support professionals
- Marketing professionals
- Product or project managers
- Line-of-business or IT managers
- Sales professionals
- Non-developer technical professionals

Prerequisites

We recommend that attendees of this course have the following:

- Familiarity with using AI tools in daily work and an understanding of basic generative AI concepts. No prior AWS experience is required.

Course outline

Module 0: Welcome and Course Overview

- Course structure: the Trust Arc and the Adoption Arc
- The four mindsets in each arc and the consumer-to-producer shift
- The two take-home deliverables: the Output Trust Card and the Governed Agent Card

Trust Arc - AI as a Companion Tool

Module 1: From Model to Companion Tool

- AI, machine learning, deep learning, and generative AI as nested ideas
- What a foundation model can and cannot do on its own
- How AWS Quick closes those gaps with connected sources, citations, and permission scope

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Generative AI Essentials on AWS

AWS Classroom Training

Module 2: Output Failures and Responsible AI

- Three AI output failure modes: hallucination, bias, and boundary violations
- Four prompt strategies to reduce failure risk: name your sources, leave out sensitive data, remove loaded framing, and ask for the right material
- Standard prompt-pattern vocabulary and the business case for responsible AI

Module 3: Retrieving Answers and Grounding Output

- Your data as a competitive advantage, and how grounding reduces hallucination
- Where AWS Quick answers come from, and why two people can get different answers
- How RAG works: retrieve, augment, generate; how documents are chunked, embedded, and indexed in a knowledge base (concept only)
- When to use RAG versus a real-time connection
- Grounded versus ungrounded output and citations as an audit trail

Lab 1: Evaluate and Govern AI Output with AWS Quick

- Evaluate, trace, fix, and stress-test AI-generated output in AWS Quick
- Ends with: Output Trust Card (a completed example that students review)

Module 4: Discussion - What Would You Automate?

- What a workflow is: trigger, retrieve, reason, draft, act
- Identify a real workflow from your job and test it against five criteria: valuable, repeatable, data-dependent, currently manual, and shareable

Adoption Arc - Build, Deploy, and Govern an Agent

Module 5: From Companion to Agentic

- What changes when AI plans, reasons, uses tools, and takes multi-step action
- Agentic AI vocabulary: tools, Model Context Protocol (MCP), and multi-agent patterns
- The orchestration trace as the record of what an agent did
- Amazon Bedrock AgentCore: agents-as-config versus agents-as-code

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Generative AI Essentials on AWS

AWS Classroom Training

Module 6: Scaling an Agent - Audience, Performance, and Cost

- Personal, team, and external as a performance-and-cost widening, not just a trust widening
- Tokens and the context window; why every query costs money at scale
- How to choose a foundation model: performance, cost, and quality
- Model customization as a cost and effort spectrum, and tool design quality

Module 7: Governing an External-Facing Agent

- Layered governance as instructions versus enforced controls: system-prompt scoping, tool restrictions via AgentCore Policy, and Amazon Bedrock Guardrails
- The threats external agents face and concrete governance techniques for each
- Transparency, explainability, and Model Cards
- Evaluating agent behavior: human-in-the-loop review and LLM-as-a-judge (concept only)

Lab 2: Building an HR Assistant with Amazon Bedrock

- Build an HR assistant grounded in company documents, connect a tool, rewrite its system prompt for all employees, and attach a guardrail
- Ends with: Governed Agent Card (a completed example that students review)

Module 8: Debrief and Wrap-Up

- Recap and key takeaways
- How to bring this back to your organization

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.