

Updated on 10/12/2026

Sign Up

Databricks GenAI Engineer Associate Certification Training

3 days (21 hours)

Overview

The Databricks GenAI Engineer Associate Certification Training prepares you to take the official Databricks Certified Generative AI Engineer Associate certification exam and to design high-performing, governed, and production-ready LLM solutions on Databricks.

This certification assesses your ability to design and implement LLM-powered applications using Databricks services. You will explore architectural choices, model selection, prompt engineering, and the processing pipelines required to address business use cases.

Our Databricks GenAI Engineer Associate Certification training will help you effectively structure your preparation for the official exam. You'll solidify your understanding of key areas of the exam syllabus, from data preparation to agent development, including RAG and semantic search.

You'll learn how to build a retrieval-augmented generation application using Vector Search, manage its lifecycle with MLflow, deploy it with Model Serving, and apply Unity Catalog's governance mechanisms.

By the end of the course, you'll be able to analyze exam scenarios, select the appropriate components, evaluate the quality of responses, and prepare for the deployment, monitoring, and securing of reliable GenAI applications.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Identify the skills and areas assessed by the Databricks Certified Generative AI Engineer Associate certification
- Design a RAG architecture tailored to business needs, quality constraints, and costs
- Prepare, index, and query content using Delta Lake, Unity Catalog, and Vector Search
- Develop, trace, and evaluate pipelines and agents with MLflow
- Deploy a GenAI application using Model Serving and implement governance and monitoring principles

Target Audience

- AI Engineer
- Machine Learning Engineer
- Data Engineer
- Data Scientist
- Python Developer

Prerequisites

- Practical experience with Python and notebook development environments
- Knowledge of the basic principles of LLMs, embeddings, and RAG
- Experience with data manipulation using SQL or PySpark
- Familiarity with API deployment and model lifecycle management concepts

Technical Requirements

- Computer running Windows 10 or 11, macOS, or Linux, with at least 8 GB of RAM
- A recent version of Google Chrome, Mozilla Firefox, or Microsoft Edge, with cookies and JavaScript enabled
- Stable internet connection, without any filtering that blocks access to the Databricks workspace
- Access to a Databricks workspace with the necessary permissions to create notebooks, catalogs, and demo resources

Curriculum for our Databricks GenAI Engineer Associate Training

[Day 1 - Morning]

Certification Framework and GenAI Architecture

- Overview of the Databricks Certified Generative AI Engineer Associate certification and its competency areas.
- Review of an exam scenario; identification of functional, technical, and business requirements.

- LLM Fundamentals: Tasks, Limitations, Context, Latency, Cost, and Selection Criteria.
- Designing a GenAI pipeline: inputs, outputs, prompts, tools, memory, and multi-step reasoning.
- Positioning of Agent Bricks for knowledge agents, multi-agent supervision, and information extraction.

[Day 1 - Afternoon]

Prompts, Data, and Augmented Retrieval

- Designing structured prompts: instructions, context, examples, output formats, and business constraints.
- Selecting document sources and filtering out content that degrades the quality of an RAG.
- Segmentation strategies: chunk size, overlap, document structure, and contextual constraints.
- Embeddings, hybrid search, metadata filters, and the role of reranking in retrieval.
- Retriever quality metrics: relevance, coverage, fidelity to the corpus, and error analysis.
- Hands-on workshop: preparing a document corpus, defining a chunking strategy, and comparing retrieval results for a business assistant.

[Day 2 - Morning]

Application and Agent Development

- Using LangChain components or equivalents in a GenAI application on Databricks.
- Selecting a generation or embedding model based on the task, quality, context window, and cost constraints.
- Building a pipeline with preprocessing, context enrichment, generation, and post-processing of responses.
- Implementing safeguards: input validation, prevention of malicious instructions, and rejection policies.
- Designing agents using tools, conversational APIs, data spaces, or MCP servers.

[Day 2 - Afternoon]

Building and deploying a RAG application

- Organizing data in Delta Lake and managing objects with Unity Catalog.
- Creating and querying a Vector Search index based on requirements for volume, freshness, latency, and cost.
- Recording a pipeline with MLflow: dependencies, signature, input examples, and model pyfunc.
- Deployment with Model Serving and using the Foundation Model APIs for inference.

- CI/CD Principles: component unit tests, prompt promotion, and controlled index updates.
- Hands-on Workshop: Assembling an RAG pipeline, registering it in MLflow, creating a Vector Search index, and prepare for deployment.

[Day 3 - Morning]

Governance, Evaluation, and Monitoring

- Using Unity Catalog to control access to data, models, endpoints, and application resources.
- Risk management: data provenance, licenses, sensitive data, masking, and prompt security.
- Evaluating a GenAI application with MLflow: test sets, metrics, LLM evaluators, and feedback from domain experts.
- Distinguishing between retrieval evaluation, generation quality, groundedness, relevance, and security.
- Production monitoring with inference tables, AI Gateway, usage tracking, cost control, and rate limiting.

[Day 3 - Afternoon]

Targeted review and exam preparation

- Review of exam topics: design, data preparation, development, deployment, governance, and monitoring.
- Analysis of multiple-choice and multiple-answer questions: keywords, implicit constraints, and elimination of distractor answers.
- Certification trade-offs: model, retrieval strategy, endpoint, security, cost, performance, and maintainability.
- Time management and proofreading techniques for the 90-minute timed exam.
- Individual preparation plan: revision priorities, practical resources, and reinforcement of key points to watch for.
- Hands-on workshop: complete a mock certification exam based on a full scenario, justify each architectural choice, and correct the answers as a group.

Target Companies

This training is intended for both individuals and companies—large or small—seeking to train their teams in a new advanced IT technology or to acquire specific professional knowledge or modern methods.

Entry-Level Assessment

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency in various types of technologies, as well as their personal expectations and objectives

regarding the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose challenges for monitoring and ensuring the smooth running of the training session.

Teaching Methods

Hands-On Course: 60% Hands-On, 40% Theory. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the instructor—supported by examples and reflection sessions—and group work.

Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

Certification

A certificate will be issued to each trainee who has completed the entire training program.