

# AWS Generative AI Developer – Professional

Course syllabus · Live online · Professional

Build production-grade generative AI applications on AWS with foundation models, RAG, agents and guardrails.

|                      |                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------|
| <b>Live training</b> | 48 hours in total                                                                         |
| <b>Weekday batch</b> | 5 weeks · Monday to Friday, 2 hours per day (10 hours a week)                             |
| <b>Weekend batch</b> | 6 weeks · Saturday and Sunday, 4 hours each day (8 hours a week)                          |
| <b>Self-practice</b> | about 40 hours in total (labs, assignments, revision)                                     |
| <b>Total effort</b>  | about 88 hours including self-practice                                                    |
| <b>Mode</b>          | Live instructor-led online sessions · session recordings after each class · small batches |
| <b>Trainer</b>       | Ameya Vaidya                                                                              |
| <b>Batch start</b>   | Shared by our counselor                                                                   |
| <b>Batch timings</b> | Weekday / weekend options, shared by our counselor                                        |
| <b>Course fee</b>    | Shared by our counselor (varies with current offers)                                      |

## Who is this for?

Developers and ML engineers building LLM-powered applications.

## Prerequisites

Python programming, AWS fundamentals and basic ML/LLM awareness. AI Practitioner or Developer Associate recommended.

## What you will learn

- Integrate foundation models into applications using Amazon Bedrock
- Design and implement RAG with vector stores
- Build agents and tool-using workflows
- Apply prompt engineering, evaluation and guardrails
- Deploy, monitor and optimize GenAI workloads for cost and latency
- Secure and govern GenAI applications

## Tools and technologies

Amazon Bedrock · Bedrock Agents · Knowledge Bases · OpenSearch · SageMaker · Lambda · Python · LangChain

## Detailed syllabus

| # | Module                                       | Hours | Topics                                                                                                                                                                          | Hands-on                  |
|---|----------------------------------------------|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| 1 | <b>GenAI Foundations and Model Selection</b> | 5     | <ul style="list-style-type: none"><li>• Foundation model landscape</li><li>• Choosing and comparing models</li><li>• Inference parameters, tokens and cost</li></ul>            | Compare models in Bedrock |
| 2 | <b>Prompt Engineering</b>                    | 5     | <ul style="list-style-type: none"><li>• Prompt patterns and structured output</li><li>• System prompts and context management</li><li>• Prompt versioning and testing</li></ul> | Build a prompt library    |

| # | Module                                         | Hours     | Topics                                                                                                                                                                                 | Hands-on                       |
|---|------------------------------------------------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| 3 | <b>Embeddings and Vector Stores</b>            | 5         | <ul style="list-style-type: none"> <li>• Embeddings and similarity search</li> <li>• OpenSearch, pgvector and Knowledge Bases</li> <li>• Chunking and indexing strategies</li> </ul>   | Index a document set           |
| 4 | <b>Retrieval-Augmented Generation</b>          | 7         | <ul style="list-style-type: none"> <li>• RAG architecture and pipelines</li> <li>• Hybrid search and re-ranking</li> <li>• Grounding and citations</li> </ul>                          | Build an enterprise Q&A bot    |
| 5 | <b>Agents and Tool Use</b>                     | 6         | <ul style="list-style-type: none"> <li>• Bedrock Agents and action groups</li> <li>• Function calling and orchestration</li> <li>• Multi-step workflows with Step Functions</li> </ul> | Agent that calls business APIs |
| 6 | <b>Customization</b>                           | 4         | <ul style="list-style-type: none"> <li>• Fine-tuning vs RAG vs prompting</li> <li>• Model customization options on AWS</li> <li>• Evaluation datasets</li> </ul>                       | -                              |
| 7 | <b>Safety, Guardrails and Governance</b>       | 5         | <ul style="list-style-type: none"> <li>• Bedrock Guardrails</li> <li>• PII protection and data privacy</li> <li>• Responsible-AI controls and auditing</li> </ul>                      | Apply guardrails to an app     |
| 8 | <b>Evaluation, Monitoring and Optimization</b> | 6         | <ul style="list-style-type: none"> <li>• Automated and human evaluation</li> <li>• Observability and tracing</li> <li>• Caching, batching and cost optimization</li> </ul>             | Evaluation harness             |
| 9 | <b>Deployment and Capstone</b>                 | 5         | <ul style="list-style-type: none"> <li>• Serverless and containerized deployment</li> <li>• CI/CD for GenAI apps</li> <li>• Capstone build and review</li> </ul>                       | Capstone project               |
|   | <b>Total live hours</b>                        | <b>48</b> |                                                                                                                                                                                        |                                |

## Projects

- Enterprise document Q&A (RAG) application
- Tool-using AI agent with guardrails

## Assessment

Weekly build assignments, evaluation review and capstone assessment.

## Certification

Prepares you for the AWS Certified Generative AI Developer – Professional exam. Exam blueprint may change; syllabus is aligned to current AWS guidance.

## What is included

- Live instructor-led online sessions
- Hands-on labs and assignments
- Session recordings after each class
- Doubt-solving support
- Course completion certificate
- Small interactive batches
- Unpaid internship opportunities

**Our promise:** we do not offer guaranteed-placement claims. We focus on building genuine skills so you are capable of clearing interviews and doing the job. Unpaid internship opportunities are available for practical exposure. Fees and batch dates vary with current offers, so please confirm them with our counselor. This syllabus is indicative and may be updated to reflect industry changes.

**Book a free demo:** <https://www.enbridgeinfotech.com/free-demo.html> · WhatsApp +91-9156091020