

Developing Generative AI Applications on AWS

AWS Classroom Training

Course description

This course is designed to introduce generative AI to software developers interested in leveraging large language models without fine-tuning. The course provides an overview of generative AI, planning a generative AI project, getting started with Amazon Bedrock, the foundations of prompt engineering, and the architecture patterns to build generative AI applications using Amazon Bedrock and LangChain.

- Course level: Advanced
- Duration: 2 days

Activities

This course includes presentations, demonstrations, and group exercises.

Course objectives

In this course, you will learn to:

- Describe generative AI and how it aligns to machine learning
- Define the importance of generative AI and explain its potential risks and benefits
- Identify business value from generative AI use cases
- Discuss the technical foundations and key terminology for generative AI
- Explain the steps for planning a generative AI project
- Identify some of the risks and mitigations when using generative AI
- Understand how Amazon Bedrock works
- Familiarize yourself with basic concepts of Amazon Bedrock
- Recognize the benefits of Amazon Bedrock
- List typical use cases for Amazon Bedrock
- Describe the typical architecture associated with an Amazon Bedrock solution
- Understand the cost structure of Amazon Bedrock
- Implement a demonstration of Amazon Bedrock in the AWS Management Console
- Define prompt engineering and apply general best practices when interacting with FMs
- Identify the basic types of prompt techniques, including zero-shot and few-shot learning
- Apply advanced prompt techniques when necessary for your use case
- Identify which prompt-techniques are best-suited for specific models
- Identify potential prompt misuses
- Analyze potential bias in FM responses and design prompts that mitigate that bias
- Identify the components of a generative AI application and how to customize a foundation model (FM)
- Describe Amazon Bedrock foundation models, inference parameters, and key Amazon Bedrock APIs
- Identify Amazon Web Services (AWS) offerings that help with monitoring, securing, and governing your Amazon Bedrock applications
- Describe how to integrate LangChain with large language models (LLMs), prompt templates, chains, chat models, text embeddings models, document loaders, retrievers, and Agents for Amazon Bedrock
- Describe architecture patterns that can be implemented with Amazon Bedrock for building generative AI applications
- Apply the concepts to build and test sample use cases that leverage the various Amazon Bedrock models, LangChain, and the Retrieval Augmented Generation (RAG) approach

Developing Generative AI Applications on AWS

AWS Classroom Training

Intended audience

This course is intended for:

- Software developers interested in leveraging large language models without fine-tuning

Prerequisites

We recommend that attendees of this course have:

- *AWS Technical Essentials*
- Intermediate-level proficiency in Python

Developing Generative AI Applications on AWS

AWS Classroom Training

Course outline

Day 1

Module 1: Introduction to Generative AI - Art of the Possible

- Overview of ML
- Basics of generative AI
- Generative AI use cases
- Generative AI in practice
- Risks and benefits

Module 2: Planning a Generative AI Project

- Generative AI fundamentals
- Generative AI in practice
- Generative AI context
- Steps in planning a generative AI project
- Risks and mitigation

Module 3: Getting Started with Amazon Bedrock

- Introduction to Amazon Bedrock
- Architecture and use cases
- How to use Amazon Bedrock
- Demonstration: Setting Up Bedrock Access and Using Playgrounds

Module 4: Foundations of Prompt Engineering

- Basics of foundation models
- Fundamentals of prompt engineering
- Basic prompt techniques
- Advanced prompt techniques
- Demonstration: Fine-Tuning a Basic Text Prompt
- Model-specific prompt techniques
- Addressing prompt misuses
- Mitigating bias
- Demonstration: Image Bias-Mitigation

Day 2

Module 5: Amazon Bedrock Application Components

- Applications and use cases
- Overview of generative AI application components
- Foundation models and the FM interface
- Working with datasets and embeddings
- Demonstration: Word Embeddings
- Additional application components
- RAG
- Model fine-tuning
- Securing generative AI applications
- Generative AI application architecture

Developing Generative AI Applications on AWS

AWS Classroom Training

Module 6: Amazon Bedrock Foundation Models

- Introduction to Amazon Bedrock foundation models
- Using Amazon Bedrock FMs for inference
- Amazon Bedrock methods
- Data protection and auditability
- Demonstration: Invoke Bedrock Model for Text Generation Using Zero-Shot Prompt

Module 7: LangChain

- Optimizing LLM performance
- Integrating AWS and LangChain
- Using models with LangChain
- Constructing prompts
- Structuring documents with indexes
- Storing and retrieving data with memory
- Using chains to sequence components
- Managing external resources with LangChain agents
- Demonstration: Bedrock with LangChain Using a Prompt that Includes Context

Module 8: Architecture Patterns

- Introduction to architecture patterns
- Text summarization
- Demonstration: Text Summarization of Small Files with Anthropic Claude
- Demonstration: Abstractive Text Summarization with Amazon Titan Using LangChain
- Question answering
- Demonstration: Using Amazon Bedrock for Question Answering
- Chatbots
- Demonstration: Conversational Interface – Chatbot with AI21 LLM
- Code generation
- Demonstration: Using Amazon Bedrock Models for Code Generation
- LangChain and agents for Amazon Bedrock
- Demonstration: Integrating Amazon Bedrock Models with LangChain Agents