Developing Generative AI Applications on AWS

Category: AI, Data & Analytics | **Vendor:** Amazon Web Services

Duration: 16.00 hours (2 days) **13.0 CPD Hours** **Rating:** ⭐ 4.6 (5,878 reviews)

Course Information

Delivery Format: Instructor Led - Online

Course Overview

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.

About This Course

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.

Who Should Attend

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

Prerequisites & Entry Requirements

General Prerequisites:

We recommend that attendees of this course have:

- Completed the Generative AI Essentials AWS instructor-led course
- Intermediate-level proficiency in Python
- Familiarity with AWS Cloud

Learning Outcomes

Upon successful completion of this course, participants will be able 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

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

Detailed Course Outline

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

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

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

, Day 1

Course Introduction

Module 1: Exploring Components of Generative AI Applications on AWS

- Understanding generative AI concepts
- Identifying AWS generative AI stack components
- Designing generative AI application components

Module 2: Programming with Amazon Bedrock

- Guiding model response generation
- Using Amazon Bedrock programmatically

Hands-on lab: Develop with Amazon Bedrock APIs

Hands-on lab: Develop Streaming Patterns with Amazon Bedrock APIs

Module 3: Applying Prompt Engineering for Developers

- Introducing prompt engineering
- Introducing prompt techniques

- Optimizing prompts for better results

Module 4: Using Amazon Bedrock APIs in Common Architectures

- Implementing architecture patterns with Amazon Bedrock APIs
- Exploring common use cases
- Adding conversational memory to extend context

Hands-on lab: Develop Conversation Patterns with Amazon Bedrock APIs[/b]

Module 5: Customizing Generative AI Responses with RAG

- Implementing Retrieval Augmented Generation (RAG)
- Using Amazon Bedrock Knowledge Bases

Hands-on lab: Develop Retrieval Augmented Generation (RAG) Applications with Amazon Bedrock Knowledge Bases[/b]

Module 6: Integrating Open Source Frameworks with Amazon Bedrock

- Invoking a foundation model in Amazon Bedrock using LangChain
- Using LangChain for context-aware responses

Hands-on lab: Develop a Generative AI Application Pattern using Open Source Frameworks and Amazon Bedrock Knowledge Bases[/b]

Day 2

Module 7: Evaluating Generative AI Application Components

- Evaluating application components
- Evaluating model output
- Evaluating RAG output
- Optimizing latency and cost

Hands-on lab: Evaluating Retrieval Augmented Generation (RAG) Applications[/b]

Module 8: Implementing Responsible AI

- Understanding responsible AI
- Mitigating bias and addressing prompt misuses
- Using Amazon Bedrock Guardrails

Hands-on lab: Securing Generative AI Applications Using Bedrock Guardrails[/b]

Module 9: Using Tools and Agents in Generative AI Applications

- Using tools
- Understanding AI agents
- Understanding open source agentic frameworks
- Understanding agent interoperability

Module 10: Developing Amazon Bedrock Agents

- Implementing Amazon Bedrock Flows
- Designing Amazon Bedrock Agents
- Developing Amazon Bedrock Inline Agents
- Designing multi-agent collaboration
- Using Amazon Bedrock AgentCore

Hands-on lab: Developing Amazon Bedrock Agents Integrated with Amazon Bedrock Knowledge Bases and Guardrails

Course Wrap-Up, Introduction to Generative AI – Art of the Possible, Planning a Generative AI Project, Getting Started with Amazon Bedrock, Foundations of Prompt Engineering, Amazon Bedrock Application Components, Amazon Bedrock Foundation Models, LangChain, Architecture Patterns

Skills You Will Learn:

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

- Develop generative AI applications using Amazon Bedrock.
- Design architecture patterns of generative AI applications.
- Configure Amazon Bedrock APIs to invoke foundation models (FMs) programmatically.
- Develop agentic AI applications by integrating Amazon Bedrock tools and open source frameworks.
- Build custom solutions with Retrieval Augmented Generation (RAG) and Amazon Bedrock Knowledge Bases.
- Integrate open source SDKs with Amazon Bedrock to build business.
- Optimize model responses by applying prompt engineering techniques.
- Evaluate generative AI application components.
- Implement responsible AI practices to protect generative AI.

Additional Course Details

Nexus Humans Developing Generative AI Applications on AWS training program is a workshop that presents an invigorating mix of sessions, lessons, and masterclasses meticulously crafted to propel your learning expedition forward.

This immersive bootcamp-style experience boasts interactive lectures, hands-on labs, and collaborative hackathons, all strategically designed to fortify fundamental concepts.

Guided by seasoned coaches, each session offers priceless insights and practical skills crucial for honing your expertise. Whether you're stepping into the realm of professional skills or a seasoned professional, this comprehensive course ensures you're equipped with the knowledge and prowess necessary for success.

While we feel this is the best course for the Developing Generative AI Applications on AWS course and one of our Top 10 we encourage you to read the course outline to make sure it is the right content for you.

Additionally, private sessions, closed classes or dedicated events are available both live online and at our training centres in Dublin and London, as well as at your offices anywhere in the UK, Ireland or across EMEA.

Frequently Asked Questions

Q: What delivery options are available for Developing Generative AI Applications on AWS?

We offer multiple delivery formats:

- Live Instructor-Led Classroom Online (Virtual/Live Online)
- Traditional Instructor-Led Classroom Training (ILT)
- On-site delivery at your offices anywhere in United Kingdom
- Private dedicated courses customized for your team

Q: How many CPD hours does this course provide?

The 2-day Developing Generative AI Applications on AWS course provides up to 13.0 CPD hours of structured learning. CPD certificates can be provided upon request.

Q: What is the duration of the Developing Generative AI Applications on AWS training?

The training takes place over 2 day(s), with each day lasting approximately 16.00 hours including breaks for lunch and refreshments.

Q: Do you provide corporate training for Developing Generative AI Applications on AWS?

Yes, we provide corporate training, dedicated training, and closed classes for Developing Generative AI Applications on AWS. Training can take place anywhere in United Kingdom including London, Manchester, Birmingham, Edinburgh, or live online allowing teams from across United Kingdom or internationally to attend.

Q: What related terms do people search for?

Popular related searches include: Developing Generative AI Applications on AWS, Machine Learning, DGAIA, Amazon Web Services, vILT

Q: Why choose Nexus Human for Developing Generative AI Applications on AWS?

Nexus Human is recognized as one of the leading training providers. Our trainers have won multiple awards including:

- Small Firms Best Trainer Award
- National Training Partner of the Year (Ireland) - Multiple Years
- Global Top 30 Instructor Awards (2012, 2019, 2021)
- Tech Excellence Award Nominations
- Learning Performance Institute (LPI) External Training Provider Sponsor 2024

Q: Are there any discount codes available?

Yes! Use discount code **PENPAL5** when booking your Developing Generative AI Applications on AWS training. Please note that only one discount code can be used per booking and cannot be combined with other special offers.

Nexus Human

Professional Training & Development

Email: info@nexushuman.com

Website: www.nexushuman.com

Phone: +353 1 889 8700 (Ireland) | +44 20 8142 8961 (UK)

© 2026 Nexus Human. All rights reserved. This brochure was generated on 14/11/2026.