Amazon Bedrock for Generative AI

This course is designed to introduce generative AI to software developers interested in using 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.

Prerequisites:

- Prior experience in working with AWS or completed AWS Cloud Practitioners training
- Intermediate level proficiency in Python

Duration: 3 days

Module 1: Introduction to Generative AI
- 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

Module 4: Foundations of Prompt Engineering
- Basics of foundation models

- Fundamentals of prompt engineering
- Basic prompt techniques
- Advanced prompt techniques
- Model-specific prompt techniques
- Addressing prompt misuses
- Mitigating bias

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
- 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

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

Module 8: Architecture Patterns
- Introduction to architecture patterns

- Text summarization
- Chatbots
- Code generation
- LangChain and agents for Amazon Bedrock