AZ-2005: Develop AI agents using Azure OpenAI and the Semantic Kernel SDK

Course Duration: 8 Hours **Course code: AZ-2005**

1. Course Overview

In this hands-on, instructor-led course, you'll build intelligent AI agents using **Azure OpenAI Service** alongside Microsoft's **Semantic Kernel SDK**. From designing agent architectures and implementing conversational flows to orchestrating tool use and ensuring secure, scalable deployments—you'll gain the end-to-end skills needed to produce production-grade AI solutions.

2. What you'll learn?

By the end of this course, you will be able to:

- Design multimodal AI agent architectures using Azure OpenAI models
- Integrate and leverage the Semantic Kernel SDK for building modular, extensible agents
- Build and manage conversational workflows, long-term memory, and tool integrations
- Implement tool-augmented reasoning: chain-of-thought, function calling, retrieval-augmented generation (RAG)
- Manage security, access controls, cost governance, and deployment considerations
- Monitor agent behavior, handle errors, and scale agents in production scenarios

3. Target Audience

This course is ideal for:

- AI developers & software engineers building intelligent conversational or task-oriented software
- Cloud solution architects looking to design agent-based AI solutions using Azure
- Technical leads who evaluate AI toolkits and frameworks for enterprise implementations
- Data scientists or ML engineers wanting to operationalize models in full agent systems

4. Pre-Requisites

Learners should have:

- Basic familiarity with Python (or C#) and REST APIs
- Experience with Azure fundamentals, including resource creation, authentication, and service configuration
- Understanding of natural language processing concepts and vector embeddings
- Exposure to OpenAI-style models, Azure OpenAI Service, or similar LLM APIs

5. Course content

Module 1: Introduction & Foundation

- Overview of conversational AI agents
- Azure OpenAI Service: model selection, endpoints, throughput, cost
- Semantic Kernel SDK: core concepts, setup, first agent

Module 2: Prompt Engineering & Memory

- Prompt design best practices
- Embeddings, retrieval-augmented generation (RAG)
- Implementing short-term and long-term memory with vectors

Module 3: Tool-Oriented Agents

- Defining and registering Semantic Functions
- Custom tools and plugins: calling external APIs, CRUD operations
- Tool orchestration and decision-making workflows

Module 4: Reasoning & Chain-of-Thought

- Using chain-of-thought prompting and function calling
- Handling multi-step reasoning and context
- Error detection and proactive mitigation strategies

Module 5: Orchestration & Agent Design Patterns

- Agent design templates: planner-executor, conversational payload
- Managing multi-agent collaborations and orchestration
- Debugging and tracing agent flows

Module 6: Security, Monitoring & Deployment

- Secure authentication: Azure AD, managed identities
- Secrets, access control, and data governance
- Logging, telemetry, cost monitoring, and scaling (APIM, Kubernetes, Functions)

Module 7: Capstone Project

- Design and implement a real-world AI agent (e.g., virtual help desk, scheduling assistant, data analyst agent)
- Integrate: tool chains, memory, model orchestration
- Production readiness: tests, deployment pipeline, monitoring dashboards

Module 8: Future Trends & Extensions

- On-premises and hybrid AI (Azure Stack, private endpoints)
- Incorporating agents in enterprise platforms (Microsoft 365, Power Platform)
- GPT-Plugins, openAPI-based agents, planning engines (AutoGPT, LangChain comparison)