AI Training on Azure AI Foundry & AI Architecture

Duration: 3 days (24 hours)

Course Overview

This intensive 3-day program is designed for AI architects, solution architects, senior developers, and platform engineers building enterprise-grade AI applications using Azure AI Foundry, Microsoft Agent Framework SDK, Foundry Knowledge (IQ), Work IQ, Azure Monitor, MCP, Copilot Studio, and Agent-to-Agent (A2A) communication patterns. The course prioritizes advanced implementation, architecture, observability, governance, security, and deployment topics. Introductory Azure AI Foundry concepts have been removed or minimized to fit the delivery timeframe.

Module 1: Advanced Prompt Engineering & Context Design

Topics

- System prompts, instruction hierarchy, and persona engineering
- Prompt chaining and decomposition strategies
- Few-shot and dynamic prompt generation patterns
- Context window optimization and token management
- Structured outputs, JSON schemas, and deterministic responses

Module 2: Azure AI Foundry SDK, Single-Agent Development & Tool Calling

Topics

- Working with models, deployments, inferencing APIs, and Foundry SDK components
- Creating and configuring a single-agent solution using Microsoft Agent Framework SDK
- Agent memory, state management, session handling, planning, and execution patterns
- Function calling, reusable tool contracts, schemas, and dynamic tool selection
- Connecting enterprise APIs, databases, and backend services
- Error handling, retries, fallback strategies, and testing production-ready agents

Module 3: RAG with Foundry Knowledge (IQ) and Work IQ

Topics

- Configuring enterprise knowledge sources and ingestion pipelines
- Using Work IQ for enterprise content retrieval
- Building Retrieval-Augmented Generation solutions
- Search relevance tuning, chunking, embeddings, and ranking optimization
- RAG optimization techniques, re-ranking, hybrid search, and context optimization

Module 4: Observability, Security, Multi-Agent Systems, MCP & A2A

Topics

- Azure Monitor integration; traces, logs, metrics, distributed tracing, and root cause analysis
- Performance debugging and optimization across agent workflows
- Content filtering, validation guardrails, policy enforcement, and prompt-injection protection

- Data privacy, compliance, governance, and responsible AI risk mitigation
- Multi-agent architecture, specialization, delegation, supervisor and peer-to-peer orchestration
- MCP architecture; integrating local and remote MCP servers, enterprise tools, security, and runtime execution
- A2A protocol, agent discovery, capability exchange, task routing, and state propagation

Module 5: Building and Deploying Agents with Microsoft Copilot Studio

Topics

- Governance, authentication, and monitoring considerations
- Testing, publishing, and deploying agents to supported channels
- Connecting knowledge sources and integrating actions using connectors and Power Automate
- Creating and configuring agents using topics, instructions, and generative answers
- Copilot Studio architecture, environments, and agent components

Module 6: Enterprise Capstone Project & Reference Architecture

Topics

- Designing an end-to-end enterprise AI solution
- Building a multi-agent application with RAG and tool calling
- Integrating MCP, Foundry IQ, Work IQ, and Azure Monitor
- Implementing guardrails, governance, and security controls
- Architecture review, performance tuning, and deployment roadmap

Hands-On Labs Included

Lab 1: Foundry SDK and Single-Agent Development

- Work with model deployments and configure the SDK environment
- Build a single agent using Microsoft Agent Framework SDK
- Implement agent memory and tool calling
- Integrate enterprise APIs and implement structured outputs

Lab 2: RAG with Foundry IQ and Work IQ

- Create a RAG pipeline using Foundry IQ
- Extend with Work IQ enterprise search
- Apply re-ranking and context optimization

Lab 3: Observability and Security

- Configure Azure Monitor tracing and analyze end-to-end agent execution
- Add guardrails and content safety controls
- Implement prompt injection protection

Lab 4: MCP and Multi-Agent Orchestration

- Connect local and remote MCP servers
- Build a multi-agent workflow
- Implement A2A communication

Lab 5: Copilot Studio

- Test and publish the agent to a supported channel

- Connect a knowledge source and configure agent responses
- Create and configure an agent in Microsoft Copilot Studio

Capstone

- Enterprise Research Assistant with:
    - Multi-Agent Architecture
    - Foundry IQ + Work IQ RAG
    - MCP Tool Integration
    - Azure Monitor Observability
    - Security Guardrails
    - A2A Collaboration