Enterprise AI with .NET and Microsoft Agent Framework(MAF)

Item	Plan
Duration	5 days
Daily effort	8hrs
Delivery mode	Live online / instructor-led
Total live hours	40hrs
Focus	MAF, RAG, Azure AI Foundry, and enterprise AI integration

Day 1: AI & LLM Foundations for Enterprise Developers

Foundations

Topics:

- Introduction to Generative AI, LLMs, and enterprise AI use cases
- Tokens, context windows, and prompt engineering techniques
- System prompts, structured prompting, and function calling concepts
- Fine-tuning vs RAG vs agentic AI decision framework
- Local LLM deployment with Ollama (Llama 3, Mistral, Phi, DeepSeek)
- Responsible AI, hallucinations, guardrails, and cost optimization

Hands-on: Set up a local LLM with Ollama and build a simple tool-calling agent demo in C#.

Day 2: Semantic Kernel Core Development

SK development

Topics:

- Semantic Kernel architecture and plugin model
- Setting up Semantic Kernel in .NET with Azure OpenAI/Ollama connectors
- Native and semantic plugins using `[KernelFunction]`
- Memory, embeddings, and vector search integration

- Planners for multi-step task decomposition
- Streaming responses and observability with LangFuse

Hands-on: Build a multi-step SK customer support assistant with plugins, memory, planner, and streaming responses.

Day 3: Microsoft Agent Framework Architecture & Core Concepts

MAF core

Topics:

- Microsoft Agent Framework 1.0 overview and architecture
- Agents, channels, adapters, service connectors, and middleware hooks
- Memory and context providers in MAF
- MCP-based dynamic tool calling
- Human-in-the-Loop (HIL) approval patterns
- DevUI debugging and monitoring of agents

Hands-on: Build a production-style MAF agent in C# with MCP tool calling, memory, and a HIL approval step.

Day 4: Multi-Agent Orchestration, RAG & Enterprise Knowledge Systems

RAG + orchestration

Topics:

- Sequential, handoff, and collaborative multi-agent patterns
- Designing agent workflows for enterprise processes
- Retrieval-Augmented Generation (RAG) architecture
- Vector databases and Azure AI Search
- Building retrieval pipelines over documents and SQL data
- Security, governance, and observability in RAG systems

Hands-on: Build a multi-agent workflow and create an end-to-end SQL-based RAG pipeline.

Day 5: Azure AI Foundry, Enterprise Integration & Capstone

Deployment + capstone

Topics:

- Azure AI Foundry: model deployment, fine-tuning, and governance
- No-code agents and Prompt Flow
- ML.NET integration for predictive intelligence

- Integrating AI agents with SQL Server, REST APIs, ERP systems, and MCP
- Monitoring, SLOs, security, and production readiness
- Capstone architecture review and demo preparation

Hands-on: Build and present an end-to-end enterprise AI solution integrating MAF, RAG, Azure AI Foundry, and enterprise APIs.

Expected Outcomes

- Understand enterprise AI architecture patterns and when to use RAG, fine-tuning, or agentic AI.
- Build AI applications using Semantic Kernel and Microsoft Agent Framework in C#.
- Implement autonomous and multi-agent workflows with MCP tool calling and memory.
- Create RAG pipelines using vector databases and Azure AI Search.
- Deploy, govern, and monitor AI solutions on Azure AI Foundry.
- Deliver a capstone enterprise AI solution demonstrating architecture, integration, and live deployment.