

Updated on 09/23/2026

Sign Up

Microsoft Foundry Training

3 days (21 hours)

Overview

Microsoft Foundry is the Microsoft platform for designing, evaluating, securing, and scaling AI applications and agents, from prototyping through to production deployment.

This Microsoft Foundry training will enable you to build generative AI solutions by selecting models tailored to your needs and leveraging the platform's services, tools, and SDKs. You'll learn to develop robust applications with Foundry Models and structure their interactions with your organization's data and services.

You'll be able to build RAG architectures to ground model responses in your business data. You'll learn how to prepare and index your data sources, improve the relevance of retrieved information, and provide contextualized and traceable responses.

The training also covers the creation of AI agents using Microsoft Foundry Agent Service. You'll learn how to connect your agents to business tools and APIs, leverage the Model Context Protocol (MCP), and implement human validation mechanisms to oversee sensitive actions.

Finally, you'll learn how to evaluate, trace, and monitor your applications and agents to identify errors, minimize hallucinations, and find the best balance between quality, latency, and cost. Security, RBAC, data protection, and governance are also integrated into the preparation for production deployment.

By the end of this course, you will be able to design and deploy an AI application or agent using Microsoft Foundry, from model selection and knowledge connection through to evaluation, security, and monitoring in production.

Like all our training courses, this one will introduce you to **the latest stable version** of the

technology and its new features.

Objectives

- Understand the architecture and key services of Microsoft Foundry
- Build AI applications using Foundry Models and the platform's SDK
- Implement a RAG architecture with search, grounding, and citations
- Evaluate, track, and optimize the quality, latency, and costs of AI applications
- Build agents using Microsoft Foundry Agent Service and connect tools via MCP
- Securing, governing, deploying, and monitoring applications and agents in production

Target Audience

- Developers and software engineers
- AI Engineers
- Data Engineers and ML Engineers
- Cloud architects and solutions architects
- Professionals responsible for scaling applications and AI agents

Prerequisites

- Practical knowledge of Python or JavaScript
- Understanding of REST APIs and the JSON format
- General knowledge of Azure: resources, identities, RBAC, and networking
- Basic understanding of generative AI, prompting, and embeddings is recommended

Technical Prerequisites

- Computer running Windows, macOS, or Linux with access to the terminal
- A modern web browser and a development environment such as Visual Studio Code
- Access to an Azure environment that allows you to use Microsoft Foundry and create the resources needed for the hands-on exercises
- Have the necessary permissions and roles to deploy models, create projects, and use the associated services
- Git and a Python or JavaScript environment configured for exercises using the Foundry SDK

Curriculum for our Microsoft Foundry training: Developing AI Applications and Agents

[Day 1 - Morning]

Introduction to Microsoft Foundry and Designing an AI Application

- Understand the evolution from Azure AI Foundry to Microsoft Foundry and its role within the Microsoft ecosystem
- Explore the new portal, resources, projects, and key components of Microsoft Foundry
- Understand the unified architecture centered around models, agents, and tools
- Identify use cases: business assistants, RAG, automation, agents, and generative AI applications
- Get started with the Foundry SDK 2.x and define constraints for data, security, quality, latency, and cost
- Hands-on workshop: Create a Microsoft Foundry project and design the architecture for your first AI application.

[Day 1 - Afternoon]

Build an application with Foundry Models

- Explore the Foundry model catalog and select a model based on use case, quality, latency, and cost
- Configure and use the endpoints required for inference from an application
- Structure prompts with instructions, context, examples, and output formats
- Leverage the Responses API and structured outputs to build robust applications
- Version prompts and set up test suites to compare model behaviors
- Hands-on workshop: Build an AI application using Foundry Models, structured JSON outputs, and non-regression tests.

[Day 2 - Morning]

Building RAG applications and grounding responses in data

- Understand the architecture of Retrieval-Augmented Generation (RAG) and its use cases in Microsoft Foundry
- Prepare document sources, metadata, and chunking strategies
- Use embeddings, indexing, and search mechanisms to retrieve relevant information
- Implement retrieval and file search capabilities to connect models to domain knowledge
- Improve grounding, citations, filtering, and traceability of generated responses
- Hands-on workshop: Build an RAG system using a document corpus and generate context-aware responses with their sources.

[Day 2 - Afternoon]

Evaluate, Map, and Optimize AI Applications

- Define quality criteria: relevance, accuracy, completeness, groundedness, and security
- Build evaluation datasets and compare multiple versions of an application
- Leverage Microsoft Foundry's evaluation capabilities to measure the quality of responses
- Implement tracing and monitoring to analyze application calls, errors, and behavior
- Optimize prompts, context, model selection, latency, and inference costs
- Hands-on workshop: Compare multiple configurations of a RAG application and select the best one based on quality, cost, and latency criteria.

[Day 3 - Morning]

Build agents with Microsoft Foundry Agent Service

- Understand the Foundry Agent Service architecture and distinguish between conversational applications, workflows, and AI agents
- Create and configure agents with instructions, context, tools, and knowledge retrieval mechanisms knowledge
- Connect agents to business tools and APIs and control the data exchanged
- Use the Model Context Protocol (MCP) to extend agent capabilities and integrations
- Implement safeguards, action restrictions, and human validation mechanisms for sensitive operations
- Hands-on workshop: Build a Foundry agent connected to a business tool and a via MCP.

[Day 3 - Afternoon]

Secure and deploy AI applications and agents

- Protect resources and data with RBAC, identities, secrets, and the principle of least privilege
- Configure the network mechanisms and controls necessary to isolate applications and agents
- Apply governance, content security, and sensitive data protection policies
- Organize versioning, deployment, configurations, and rollback strategies
- Monitor applications and agents in production and use logs, traces, metrics, and feedback to improve their reliability
- Hands-on workshop: Build a production deployment plan that integrates security, governance, deployment, monitoring, and quality metrics.

Target Audience

This training is intended for both individuals and companies—large or small—that wish to train their teams in a new advanced IT technology or to acquire specific business knowledge or modern methodologies.

Assessment upon enrollment

The pre-training assessment complies with Qualiopi quality standards. Upon

final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency with various types of technologies, as well as their expectations and personal goals for the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose challenges for the monitoring and smooth operation of the training session.

Teaching Methods

Hands-On Training: 60% Practical, 40% Theoretical. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the instructor—supported by examples and reflection sessions—and group work.

Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

Certification

A certificate will be issued to each trainee who has completed the entire training program.