

Updated on September 11, 2026

[Sign Up](#)

Azure AI Apps and Agents Developer Certification Training (AI-103)

Formerly AI-102

4 days (28 hours)

Overview

The Microsoft Certified Azure AI Apps and Agents Developer Associate (AI-103) certifies your ability to design, develop, secure, and deploy generative AI applications and agents using Microsoft Foundry and Azure AI services.

Our AI-103 training will enable you to design GenAI applications and AI agents capable of leveraging models, tools, data, and external services in a Microsoft Azure environment.

You'll learn how to use Microsoft Foundry to select and deploy models, develop applications with Microsoft SDKs, design effective prompts, and implement RAG architectures with Azure AI Search. You'll also explore grounding, vector search, and data enrichment mechanisms.

This training will enable you to build AI agents with memory, tools, and custom functions, and then orchestrate multi-agent scenarios. You'll learn how to connect these agents to APIs, knowledge bases, and business services while integrating validation and control mechanisms.

You will also develop solutions for natural language processing, speech recognition, computer vision, and information extraction. Multimodal capabilities and Azure Content Understanding will enable you to process documents, images, audio files, and videos in application scenarios.

Finally, you'll learn how to evaluate, secure, and monitor your solutions using observability mechanisms, content filters, safeguards, identity management, and Responsible AI practices. You'll also cover deployment, CI/CD, and

preparation for the AI-103 certification.

Upon completion of this training, you will be able to develop and scale AI applications and agents with Microsoft Foundry, from design and grounding through deployment, security, and monitoring.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Design and manage AI solutions with Microsoft Foundry
- Develop generative AI applications connected to Azure models and services
- Implement RAG architectures and grounding mechanisms with Azure AI Search
- Create and orchestrate AI agents using memory, tools, APIs, and custom functions
- Develop solutions for text processing, speech recognition, computer vision, and information extraction
- Evaluate, secure, monitor, and deploy AI applications and agents in production
- Prepare for the Azure AI Apps and Agents Developer Associate (AI-103) certification exam

Target Audience

- Python developers
- Cloud developers
- AI Engineers
- Azure Engineers
- AI Solution Architects
- Professionals interested in developing AI applications and agents with Microsoft Azure

Prerequisites

- Have hands-on experience with Python development
- Familiarity with the fundamentals of Microsoft Azure and cloud architectures
- Understand the principles of LLMs, generative AI, and prompt engineering
- Be familiar with using APIs and SDKs in an application

Technical Prerequisites

- Computer running Windows 10 or 11, macOS, or Linux, with at least 8 GB of RAM
- A Microsoft account linked to an Azure subscription that allows access to the AI services required for the workshops
- Python environment with the Microsoft Foundry and Azure SDKs required for development and exercises

- Azure AI Search and Content Understanding resources prepared for RAG and information extraction exercises

Curriculum for our Azure AI Apps and Agents Developer (AI-103) certification training

[Day 1 - Morning]

Design and manage an AI solution with Microsoft Foundry

- Understand the role of Microsoft Foundry in building AI applications and agents on Azure
- Choose the right models for different use cases: LLMs, lightweight models, multimodal models, and code models
- Select the appropriate services for generation, grounding, vector search, and agent-based workflows
- Design the Azure infrastructure required to deploy AI applications and agents
- Configure model deployments, quotas, throughput limits, scalability, and cost management
- Implement security principles, managed identity, private networking, and access control

[Day 1 - Afternoon]

Develop generative AI applications

- Deploy and use generative models with Microsoft Foundry SDKs
- Design prompts and adjust generation parameters based on application requirements
- Build multi-step workflows combining models, rules, and external tools
- Integrate a Python application into a Microsoft Foundry project
- Evaluate the relevance, quality, safety, and hallucinations produced by a GenAI application
- Hands-on workshop: Develop your first GenAI application connected to Microsoft Foundry and analyze the quality of its responses

[Day 2 - Morning]

Implementing RAG and grounding pipelines

- Understanding the architecture of a Retrieval-Augmented Generation (RAG) solution
- Ingest and index content from documents, images, audio, and video files
- Configure semantic, hybrid, and vector search to improve grounding
- Implement the enrichment, chunking, and extraction mechanisms required for indexing
- Connect search pipelines to the applications and tools used by agents
- Monitor the quality of the data ingestion, the status of the indexes, and the relevance of the results

[Day 2 - Afternoon]

Building AI agents with Microsoft Foundry

- Define the role, objectives, and behavior of an AI agent
- Implement function calling, tools, and conversation memory mechanisms
- Connect an agent to APIs, knowledge bases, search tools, and custom functions
- Build autonomous or semi-autonomous workflows with validations and safeguards
- Orchestrating multiple specialized agents in a multi-agent scenario
- Hands-on workshop: Enhance the Day 1 application with an agent equipped with memory, tools, and access to a knowledge base

[Day 3 - Morning]

Evaluate, secure, and monitor AI applications and agents

- Implement metrics to measure the quality, relevance, and security of responses
- Configure guardrails, content filters, and risk detection mechanisms
- Track interactions with models and agents using observability mechanisms
- Analyze tokens, latency, errors, agent behavior, and grounding quality
- Implement auditing, traceability, and approval workflows
- Apply Responsible AI principles to generative and agent-based solutions

[Day 3 - Afternoon]

Develop natural language and speech processing solutions

- Extract entities, topics, summaries, and structured data from textual content
- Analyze sentiment, tone, sensitive content, and security signals
- Implement translation scenarios using Microsoft Foundry services
- Create speech-to-text and text-to-speech workflows
- Integrate speech as a modality into an application or AI agent
- Hands-on workshop: Add text or speech analysis capabilities to the overarching project

[Day 4 - Morning]

Develop vision and multimodal AI solutions

- Analyze images using multimodal models and interpret their context
- Generating descriptions, captions, and responses based on visual elements
- Implement image generation and editing using prompts
- Process and analyze video sequences with Microsoft Foundry tools
- Use Azure Content Understanding to extract features from visual content

- Implement security and moderation policies tailored to multimodal content

[Day 4 - Afternoon]

Extract information, streamline processes, and prepare for the AI-103 exam

- Build information extraction workflows from documents, images, audio, and video
- Combine OCR, search, enrichment, and Content Understanding in an application pipeline
- Integrate Microsoft Foundry projects into a CI/CD pipeline
- Prepare for the deployment, monitoring, and maintenance of models and agents in production
- Review the official exam topics and weightings for the AI-103 exam
- Final workshop: finalize and secure the agent-based application developed during the training, then analyze a scenario

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.

Entry-Level Assessment

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 in 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 monitoring and ensuring the smooth running 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.

