

AI-103T00: DEVELOP AI APPS AND AGENTS ON AZURE



DURATION	LEVEL	TECHNOLOGY	DELIVERY METHOD	TRAINING CREDITS
4 Days	Intermediate	Azure / Microsoft Foundry	Instructor-led / VILT	N/A

INTRODUCTION

This course is intended for software developers wanting to build AI infused applications that leverage Microsoft Foundry. Topics in this course include developing generative AI apps, building AI agents, and solutions that implement knowledge connections or tools in your agentic applications. This course also covers multimodal capabilities and understanding of complex content.

AUDIENCE PROFILE

This course was designed for software engineers concerned with building, managing and deploying AI solutions that leverage Microsoft Foundry. They are familiar with Python and have knowledge on using APIs and SDKs to build agents and generative AI solutions on Azure.

PREREQUISITES

- Basic software development concepts
- Be familiar with Azure services and the Azure portal.
- Be familiar with generative AI concepts and use cases.
- Have experience with Python or another programming language.
- Basic familiarity with Visual Studio Code
- It's highly recommended you have experience with the Microsoft Foundry Agent Service.
- A Microsoft 365 subscription with access to Teams

COURSE OBJECTIVES

After completing this course, students will be able to:

- Should be familiar with fundamental AI concepts and services in Azure. You should also have programming experience.
- Explore and filter models in the model catalog
- Compare models using benchmark metrics for quality, safety, cost, and performance
- Deploy a model to an endpoint and test it in the playground

COURSE CONTENT

Module 1: Plan and prepare to develop AI solutions on Azure

Microsoft Azure offers multiple services that enable developers to build amazing AI-powered solutions. Proper planning and preparation involves identifying the services you'll use and creating an optimal working environment for your development team.

Lessons:

- Introduction
- What is AI?
- Microsoft Foundry
- Foundry Tools
- Developer tools and SDKs
- Responsible AI
- Exercise - Prepare for an AI development project

- Module assessment
- Summary

After completing this module, students will be able to:

- Identify common AI capabilities that you can implement in applications
- Describe Microsoft Foundry and considerations for using it
- Describe Foundry Tools and considerations for using them
- Identify appropriate developer tools and SDKs for an AI project
- Describe considerations for responsible AI

Module 2: Select, deploy, and evaluate Microsoft Foundry models

Explore how to select appropriate models from the model catalog using benchmarks, deploy them to endpoints, and evaluate their performance using manual and automated approaches in Microsoft Foundry portal.

Lessons:

- Introduction
- Explore the model catalog
- Select models using benchmarks
- Deploy models to endpoints
- Evaluate model performance
- Exercise - Select, deploy, and evaluate models

- Knowledge check
- Summary

After completing this module, students will be able to:

- Explore and filter models in the model catalog
- Compare models using benchmark metrics for quality, safety, cost, and performance
- Deploy a model to an endpoint and test it in the playground
- Evaluate model performance using manual and automated approaches
- Understand different evaluation metrics and when to use them

Module 3: Develop a generative AI chat app with Microsoft Foundry

Use Microsoft Foundry to develop generative AI chat applications with projects and the Responses API.

Lessons:

- Introduction
- Explore with the model playground
- Choose an endpoint and SDK
- Generate responses with the Responses API
- Generate responses with the ChatCompletions API
- Exercise - Create a generative AI chat app
- Knowledge check
- Summary

After completing this module, students will be able to:

- Describe the process for creating a generative AI chat application with Microsoft Foundry.
- Use the Chat playground to explore models and generate code samples.
- Choose an endpoint, authentication method, and client SDK for your app development.
- Use the Responses API to generate AI responses in applications.
- Use the ChatCompletions API to generate AI responses in applications.

Module 4: Develop generative AI apps that use tools

Tools enable models to perform tasks and interact with external systems, enabling them to extend their capabilities beyond basic chat interactions.

Lessons:

- Introduction
- What are tools?
- Use the code_interpreter tool
- Use the web_search tool
- Use the file_search tool
- Use the function tool

- Exercise - Create a generative AI chat app that uses tools
- Module assessment
- Summary

After completing this module, students will be able to:

- Describe the capabilities of generative AI tools.
- Use the code_interpreter tool to run code and analyze data.
- Use the web_search tool to retrieve real-time information from the internet.
- Use the file_search tool to access and analyze files.
- Use the function tool to run custom code.

Module 5: Optimize generative AI model performance with Microsoft Foundry

Explore complementary strategies to optimize generative AI model performance. Learn how to apply prompt engineering, ground your model with RAG, and fine-tune for consistent behavior—and when to combine these approaches.

Lessons:

- Introduction
- Optimize model output with prompt engineering
- Ground your model with Retrieval Augmented Generation
- Fine-tune a model for consistent behavior
- Compare and combine optimization strategies
- Exercise - Optimize generative AI model performance
- Module assessment
- Summary

After completing this module, you will be able to:

- Apply prompt engineering techniques including system messages, few-shot learning, and model parameters to optimize model output.
- Understand when and how to ground a language model using Retrieval Augmented Generation (RAG).
- Identify when fine-tuning a model improves behavioral consistency.
- Compare optimization strategies and determine when to combine them.

Module 6: Implement a responsible generative AI solution in Microsoft Foundry

Generative AI enables amazing creative solutions but must be implemented responsibly to minimize the risk of harmful content generation.

Lessons

- Introduction
- Plan a responsible generative AI solution

- Map potential harms
- Measure potential harms
- Mitigate potential harms
- Manage a responsible generative AI solution
- Exercise - Apply guardrails to prevent the output of harmful content
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe an overall process for responsible generative AI solution development
- Identify and prioritize potential harms relevant to a generative AI solution
- Measure the presence of harms in a generative AI solution
- Mitigate harms in a generative AI solution
- Prepare to deploy and operate a generative AI solution responsibly

Module 7: Develop AI agents with Microsoft Foundry and Visual Studio Code

Learn how to build, test, and deploy AI agents using Microsoft Foundry Agent Service through both the Azure portal and Visual Studio Code extension.

Lessons

- Introduction
- Understand AI agents and Microsoft Foundry Agent Service
- Explore development approaches
- Build your first agent in Microsoft Foundry
- Set up Visual Studio Code for agent development
- Configure and manage agents in Visual Studio Code
- Extend agent capabilities with tools
- Test, deploy, and integrate agents
- Exercise - Build and deploy an AI agent
- Knowledge check
- Summary

After completing this module, you will be able to:

- Describe the purpose and capabilities of AI agents
- Explain the key features of Microsoft Foundry Agent Service
- Set up and configure the Microsoft Foundry extension in Visual Studio Code
- Build and configure AI agents using multiple development approaches
- Extend agent capabilities with tools and functions
- Test agents using integrated playgrounds

- Deploy and integrate agents into applications

Module 8: Integrate custom tools into your agent

Built-in tools are useful, but they may not meet all your needs. In this module, learn how to extend the capabilities of your agent by integrating custom tools for your agent to use.

Lessons

- Introduction
- Why use custom tools
- Options for implementing custom tools
- How to integrate custom tools
- Exercise - Build an agent with custom tools
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe the benefits of using custom tools with your agent.
- Explore the different options for custom tools.
- Build an agent that integrates custom tools using the Microsoft Foundry Agent Service.

Module 9: Integrate MCP Tools with Azure AI Agents

Enable dynamic tool access for your Azure AI agents. Learn how to connect MCP-hosted tools and integrate them seamlessly into agent workflows.

Lessons

- Introduction
- Understand MCP tool discovery
- Integrate agent tools using an MCP server and client
- Use Azure AI agents with MCP servers
- Exercise - Connect MCP tools to Azure AI Agents
- Module assessment
- Summary

After completing this module, you will be able to:

- Explain the roles of the MCP server and client in tool discovery and invocation.
- Wrap MCP tools as asynchronous functions and register them with Azure AI agents.
- Build an Azure AI agent that dynamically accesses and calls MCP tools during runtime.

Module 10: Build knowledge-enhanced AI agents with Foundry IQ

Learn how to connect AI agents with enterprise knowledge using Foundry IQ. You'll explore how Retrieval Augmented Generation

(RAG) solves the knowledge problem for AI agents, discover how Foundry IQ provides a shared knowledge platform that multiple agents can access, improve retrieval quality through data optimization, and configure agent instructions for consistent, cited responses.

Lessons

- Introduction
- Understanding RAG for agents
- Explore Foundry IQ
- Configure data sources for knowledge bases
- Configure retrieval with Foundry IQ
- Exercise - Integrate an AI agent with Foundry IQ
- Knowledge check
- Summary

After completing this module, you will be able to:

- Explain how RAG solves the knowledge problem by connecting agents to real-time information
- Describe how Foundry IQ provides a shared knowledge platform that multiple agents can access
- Configure data sources for knowledge bases including Azure AI Search, Blob Storage, SharePoint, and OneLake
- Configure agent instructions to control retrieval behavior and ensure consistent citations
- Test and monitor agent retrieval to maintain quality in production

Module 11: Integrate your agent with Microsoft 365

Learn how to publish Microsoft Foundry agents to Microsoft Teams and Microsoft 365 Copilot, access workplace data with Work IQ, and test your integrated agents.

Lessons

- Introduction
- Understand Foundry agent publishing options
- Publish an agent from Foundry portal to Teams
- Advanced - Use Microsoft 365 Agents Toolkit
- Access Microsoft 365 data with Work IQ
- Test and iterate your integrated agent
- Exercise - Publish a Foundry agent to Teams
- Knowledge check
- Summary

After completing this module, you will be able to:

- Explain the options for publishing Foundry agents to Microsoft 365

- Publish an agent from the Foundry portal to Teams and Microsoft 365 Copilot
- Use Work IQ to access Microsoft 365 data in your agents
- Test and troubleshoot agents integrated with Microsoft 365

Module 12: Build agent-driven workflows using Microsoft Foundry

Workflows enable you to orchestrate AI agents and other components to create intelligent applications. Learn how to build and manage workflows using Microsoft Foundry.

Lessons

- Introduction
- Understand Workflows
- Identify Workflow Patterns
- Create workflows in Microsoft Foundry
- Add Agents to a Workflow
- Apply Power Fx in Workflows
- Maintain Workflows in Microsoft Foundry
- Use workflows in code
- Exercise - Create an Agent-driven Workflow
- Module Assessment
- Summary

After completing this module, you will be able to:

- Explain how nodes, variables, and agent outputs control workflow execution
- Route requests using structured outputs and conditional logic
- Loop over multiple inputs with For-Each nodes
- Use human-in-the-loop and escalation patterns for low-confidence items
- Utilize Power Fx expressions to manipulate data and control flow.

Module 13: Develop an AI agent with Microsoft Agent Framework

This module provides engineers with the skills to begin building Microsoft Foundry Agent Service agents with Microsoft Agent Framework.

Lessons

- Introduction
- Understand Microsoft Agent Framework AI agents
- Create an Azure AI agent with Microsoft Agent Framework
- Add tools to Azure AI agent
- Exercise - Develop an Azure AI agent with the Microsoft Agent Framework SDK
- Knowledge check
- Summary

After completing this module, you will be able to:

- Use Microsoft Agent Framework to connect to a Microsoft Foundry project
- Create Microsoft Foundry Agent Service agents using the Microsoft Agent Framework SDK
- Integrate plugin functions with your AI agent

Module 14: Orchestrate a multi-agent solution using the Microsoft Agent Framework

Learn how to use the Microsoft Agent Framework SDK to develop your own AI agents that can collaborate for a multi-agent solution.

Lessons

- Introduction
- Understand the Microsoft Agent Framework
- Understand agent orchestration
- Use concurrent orchestration
- Use sequential orchestration
- Use group chat orchestration
- Use handoff orchestration
- Use Magentic orchestration
- Exercise - Develop a multi-agent solution
- Knowledge check
- Summary

After completing this module, you will be able to:

- Build AI agents using the Microsoft Agent Framework SDK
- Understand how and when to use different orchestration patterns
- Develop multi-agent solutions

Module 15: Discover Azure AI Agents with A2A

Learn how to implement the A2A protocol to enable agent discovery, direct communication, and coordinated task execution across remote agents.

Lessons

- Introduction
- Define an A2A agent
- Implement an agent executor
- Host an A2A server
- Connect to your A2A agent
- Exercise - Connect to remote Azure AI Agents with the A2A protocol
- Module assessment
- Summary

After completing this module, you will be able to:

- Understand the A2A protocol and its role in multi-agent orchestration.
- Design discoverable agents for modular, collaborative problem-solving.
- Implement A2A strategies to discover and invoke remote agents.

Module 16: Analyze text with Azure Language in Foundry Tools

Azure Language in Foundry Tools enables you to create intelligent apps and services that extract semantic information from text.

Lessons

- Introduction
- Azure Language in Microsoft Foundry Tools
- Detect language
- Extract entities
- Extract personally identifiable information (PII)
- Exercise - Analyze text
- Module assessment
- Summary

After completing this module, you will be able to:

- Detect language from text.
- Recognize named entities in text.
- Extract personally identifiable information (PII) in text.

Module 17: Develop a text analysis agent with the Azure Language MCP server

Learn how to build an AI agent that uses the Azure Language MCP server to perform text analysis tasks like language detection, entity recognition, and personal information redaction.

Lessons

- Introduction
- Understand the Azure Language MCP server
- Connect and use the Language MCP server with an agent
- Exercise - Develop a text analysis agent
- Knowledge check
- Summary

After completing this module, you will be able to:

- Describe the Azure Language MCP server and the text analysis capabilities it exposes.
- Explain how MCP enables dynamic tool discovery and selection by AI agents.
- Connect the Azure Language MCP server to an agent in Microsoft Foundry.
- Build a Python client application that invokes an agent to perform text analysis.

Module 18: Develop a speech-capable generative AI application

A voice carries meaning beyond words. Learn how to use models that transcribe and synthesize speech.

Lessons

- Introduction

- Choose a speech-capable model
- Transcribe speech
- Synthesize speech
- Exercise - Use speech-capable generative AI models
- Module assessment
- Summary

After completing this module, you will be able to:

- Deploy speech-capable generative AI models in Microsoft Foundry.
- Use a generative AI model to transcribe speech.
- Use a generative AI model to synthesize speech.

Module 19: Create speech-enabled apps with Azure Speech in Microsoft Foundry Tools

Azure Speech in Microsoft Foundry Tools enables you to build speech-enabled applications. This module focuses on using the speech-to-text and text to speech APIs, which enable you to create apps that are capable of speech recognition and speech synthesis.

Lessons

- Introduction
- Azure Speech in Foundry Tools
- Use the Speech to Text API
- Use the Text to Speech API
- Configure audio format and voices
- Use Speech Synthesis Markup Language
- Exercise - Create a speech-enabled app
- Module assessment
- Summary

After completing this module, you will be able to:

- Use a Microsoft Foundry resource for Azure Speech in Foundry Tools
- Implement speech recognition with the Azure Speech to text API
- Use the Text to speech API to implement speech synthesis
- Configure audio format and voices
- Use Speech Synthesis Markup Language (SSML)

Module 20: Develop a speech agent with the Azure Speech MCP server

Learn how to build an AI agent that uses the Azure Speech MCP server to perform speech-to-text and text-to-speech tasks.

Lessons

- Introduction
- Understand the Azure Speech MCP server
- Connect and use the Speech MCP server with an agent

- Exercise - Use Azure Speech in an agent
- Knowledge check
- Summary

After completing this module, you will be able to:

- Describe the Azure Speech MCP server and the speech capabilities it exposes.
- Explain how MCP enables dynamic tool discovery and selection by AI agents.
- Set up Azure Blob Storage for audio file input and output.
- Connect the Azure Speech MCP server to an agent in Microsoft Foundry.
- Build a Python client application that invokes an agent to perform speech tasks.

Module 21: Develop an Azure Speech Voice Live Agent in Microsoft Foundry

Learn how to develop a Voice Live agent using the Voice Live API and SDK. This module covers the fundamentals of the Voice Live platform, including API integration, SDK usage, and building conversational AI agents.

Lessons

- Introduction
- Explore the Azure Voice Live API
- Explore the AI Voice Live client library for Python
- Create a Voice Live agent
- Exercise - Develop a Voice Live agent
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe the core components and capabilities of the Azure Speech Voice Live platform.
- Use the Voice Live API to create conversational AI solutions.
- Use the Voice Live SDK to build and deploy conversational AI solutions.
- Integrate Microsoft Foundry agents with the Voice Live API to create voice agents.

Module 22: Translate text and speech with Microsoft Foundry Tools

The Translator and Speech services enable you to create intelligent apps and services that can translate text and speech between languages.

Lessons

- Introduction
- Translation in Microsoft Foundry
- Translate text
- Translate speech

- Exercise - Translate text and speech
- Module assessment
- Summary

After completing this module, you will be able to:

- Identify options for translating text and speech in Microsoft Foundry Tools
- Use Azure Translator in Foundry Tools to translate text
- Use Azure Speech in Foundry Tools to translate speech

Module 23: Develop a vision-enabled generative AI application

A picture says a thousand words, and multimodal generative AI models can interpret images to respond to visual prompts. Learn how to build vision-enabled chat apps.

Lessons

- Introduction
- Use a vision-capable model in the Microsoft Foundry portal
- Develop a vision-based chat app
- Exercise - Develop a vision-enabled chat app
- Module assessment
- Summary

After completing this module, you will be able to:

- Deploy a vision-enabled generative AI model in Microsoft Foundry.
- Test an image-based prompt in the chat playground.
- Create a chat app that submits image-based prompts.

Module 24: Generate images with AI

In Microsoft Foundry, you can use image generation models to create original images based on natural language prompts.

Lessons

- Introduction
- What are image-generation models?
- Explore image-generation models in Microsoft Foundry portal
- Create a client application that uses an image generation model
- Exercise - Generate images with AI
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe the capabilities of image generation models
- Use the Images playground in Microsoft Foundry portal
- Integrate image generation models into your apps

Module 25: Generate videos with Microsoft Foundry

Learn how to generate videos from text prompts with Sora 2 in Microsoft Foundry.

Lessons

- Introduction
- Deploy a video generating model
- Generate video from a prompt
- Generate video in Python
- Exercise - Generate video with Sora 2 in Microsoft Foundry
- Module assessment
- Summary

After completing this module, you will be able to:

- Deploy a video-generating AI model in Microsoft Foundry.
- Test a video-based prompt in the chat playground.
- Use the Azure OpenAI SDK to generate and remix videos in Python.

Module 26: Analyze images with Content Understanding

Learn how to analyze images with Azure Content Understanding.

Lessons

- Introduction
- What is Content Understanding?
- Analyze images with Content Understanding
- Exercise - Analyze images with Content Understanding
- Module assessment
- Summary

After completing this module, you will be able to:

- Deploy a content-understanding AI model in Microsoft Foundry.
- Test an image-based prompt in the chat playground.
- Use the Azure OpenAI SDK to analyze images in Python.

Module 27: Create a multimodal analysis solution with Azure Content Understanding

Use Azure Content Understanding for multimodal content analysis and information extraction.

Lessons

- Introduction
- What is Azure Content Understanding?
- Create a Content Understanding analyzer
- Use the Content Understanding API
- Exercise - Extract information from multimodal content
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe capabilities of Azure Content Understanding.

- Use Azure Content Understanding to build a content analyzer.
- Consume a Content Understanding analyzer by using the API.

Module 28: Create an Azure Content Understanding client application

Use the Azure Content Understanding API for multimodal content analysis and information extraction.

Lessons

- Introduction
- Prepare to use the AI Content Understanding API
- Create a Content Understanding analyzer
- Analyze content
- Exercise - Develop a Content Understanding client application
- Module assessment
- Summary

After completing this module, you will be able to:

- Use the Azure Content Understanding API to build a content analyzer.
- Use the Azure Content Understanding API to consume an analyzer.

Module 29: Extract data with Azure Document Intelligence

Azure Document Intelligence uses OCR and deep learning models to extract text, key-value pairs, tables, and structured data from forms and documents. Learn how to use prebuilt and custom models to automate document processing.

Lessons

- Introduction
- What is Azure Document Intelligence?
- Use the Document Intelligence Studio
- Use prebuilt models
- Train and use custom models
- Exercise - Analyze documents with Document Intelligence
- Module assessment
- Summary

After completing this module, you will be able to:

- Describe the capabilities and components of Azure Document Intelligence.
- Use the Document Intelligence Studio to explore and test models.
- Use prebuilt models to extract data from common document types.
- Train and use custom models for industry-specific forms.

Module 30: Create a knowledge mining solution with Azure AI Search

Unlock the hidden insights in your data with Azure AI Search. In this module, you'll learn how to implement a knowledge mining solution that extracts and enriches data, making it searchable and ready for deeper analysis.

Lessons

- Introduction
- What is Azure AI Search?
- Extract data with an indexer
- Enrich extracted data with AI skills
- Search an index
- Persist extracted information in a knowledge store
- Exercise - Create a knowledge mining solution
- Module assessment
- Summary

After completing this module, you will be able to:

- Implement indexing with Azure AI Search
- Use AI skills to enrich data in an index
- Search an index to find relevant information
- Persist extracted information in a knowledge store.

ASSOCIATED CERTIFICATIONS & EXAM

This course will prepare delegates to write the Microsoft AI-103: Develop AI apps and agents on Azure Exam.