

Microsoft AI-103 Study Tracker

Exam AI-103: Developing AI Apps and Agents on Azure | studyfromnotes.com

Exam facts

Credential	Microsoft Certified: Azure AI Apps and Agents Developer Associate
Replaces	AI-102, retired 30 June 2026
Price	US \$165 in the United States; varies by country of proctoring
Duration	Most prep guides list 120 minutes; Microsoft does not publish the duration on the certification page, so check your booking confirmation
Passing score	700 or greater (Microsoft scaled scoring)
Domains	5 weighted domains - see page 2 checklist
Renewal	Annual, by free online assessment on Microsoft Learn

My 6-week plan - tick each week off

■ Week 1 - Foundations

Foundry tour, model catalog, Python SDK set up, tokens, prompts, embeddings. Build: deploy one model and call it from a Python script.

■ Week 2 - Plan and manage (25-30%)

Model and service selection, deployment options, CI/CD, quotas and cost, security, responsible AI controls. Build: secure a Foundry project with managed identity and keyless credentials.

■ Week 3 - Generative AI core

RAG design, indexing, retrieval methods, grounding quality, evaluation. Build: a RAG app over your own documents.

■ Week 4 - Agents (30-35% domain)

Agent roles, tool schemas, function calling, memory, multi-agent orchestration, tracing, safeguards. Build: a two-agent workflow with an approval step.

■ Week 5 - Vision, text, extraction (10-15% each)

Multimodal workflows, speech, translation, OCR and Content Understanding pipelines. Build: a document pipeline from PDF to structured JSON. Short on time? Split this week across two weeks, vision first, then text and extraction.

■ Week 6 - Revision and rehearsal

Official practice assessment on AI Skills Navigator, exam sandbox, weak-domain drills, two timed runs. Then book the exam.

My exam date: _____ My score target: _____ Practice scores: _____

Book with a personal Microsoft account, not a work account: Microsoft warns that records tied to an employer account are lost when you leave.

AI-103 Domain Checklist

Every topic below comes from the official Microsoft skills outline. Tick each one as you master it.

1. Plan and manage an Azure AI solution (25-30%)

- Choose the right model for a task: LLM, small model, multimodal model, Foundry Tools
- Choose Foundry services for generation, grounding, vector search, agent workflows
- Choose retrieval and indexing methods; memory, tool and knowledge integration
- Design Azure infrastructure and deployment options for AI apps and agents
- Configure model and agent deployments; integrate Foundry projects with CI/CD
- Manage quotas, scaling, rate limits and cost for model and agent workloads
- Monitor performance, drift, safety events, grounding and index health
- Configure security: managed identity, private networking, keyless credentials, roles
- Responsible AI: safety filters, guardrails, evaluators, trace logging, agent oversight

2. Implement generative AI and agentic solutions (30-35%)

- Deploy and consume LLMs, small models, code models and multimodal models
- Implement retrieval-augmented generation (RAG) in an application
- Design workflows, tool-augmented flows and multistep reasoning pipelines
- Evaluate models and apps: fabrications, relevance, quality and safety
- Integrate generative workflows with Foundry SDKs and connectors
- Define agent roles, goals, conversation tracking and tool schemas
- Build agents with retrieval, function calling and conversation memory
- Integrate agent tools: APIs, knowledge stores, search, custom functions
- Implement orchestrated multi-agent solutions with safeguards and approval controls
- Optimize: prompt tuning, reflection and self-critique loops, tracing, token analytics

3. Implement computer vision solutions (10-15%)

- Generate images and videos from text prompts and reference media
- Configure image editing: inpainting, mask-based edits, prompt-driven changes
- Build multimodal understanding: captions, visual question answering, alt-text
- Configure Azure Content Understanding pipelines for visual content
- Analyze video segments and identify objects, components and regions
- Apply visual safety: content filters, indirect prompt injection in images, policy rules

4. Implement text analysis solutions (10-15%)

- Extract entities, topics, summaries and structured JSON with generative prompting
- Detect sentiment, tone, safety issues and sensitive content
- Translate text with Azure Translator in Foundry Tools or LLM flows
- Customize outputs for domain tasks such as compliance summarization
- Implement speech to text and text to speech for agentic interactions
- Build retrieval and grounding pipelines: semantic, hybrid and vector search, OCR ingestion

5. Implement information extraction solutions (10-15%)

- Combine OCR, layout analysis and field extraction in multimodal pipelines
- Produce clean, grounded representations for agents and RAG with Content Understanding
- Build analyzers that output structured or markdown results for downstream reasoning