6-Month Learning Roadmap: Azure Infrastructure to AI with Microsoft Foundry

6-month mentor roadmap based on Microsoft Learn

Imoh Etuk

Microsoft Azure MVP

Learner Profile

- Background: Azure cloud infrastructure
- Target outcome: extend into AI while strengthening Azure infrastructure skills for deploying and operating AI solutions with Microsoft Foundry
- Plan duration: 6 months
- Suggested weekly cadence: 6 to 8 hours per week
- Roadmap window: May 18, 2026 to November 15, 2026

Roadmap Goals

By the end of this plan, the mentee should be able to:

1. Explain core AI, generative AI, agent, and RAG concepts in practical Azure terms.
2. Use Microsoft Foundry to select, deploy, evaluate, and govern models.
3. Design Azure infrastructure for AI solutions with secure networking, identity, monitoring, and cost controls.
4. Build a small production-style AI solution using Azure services around Microsoft Foundry.
5. Apply DevOps and MLOps practices to automate deployment and lifecycle management.

Microsoft Learn Sources Used

All learning paths below are from Microsoft Learn:

- Introduction to AI in Azure
- Get started with AI applications and agents on Azure
- Create custom copilots with Azure AI Foundry
- Manage AI-ready infrastructure
- Introduction to Cloud Infrastructure: Describe Azure architecture and services
- Introduction to Cloud Infrastructure: Describe Azure management and governance
- Build great solutions with the Microsoft Azure Well-Architected Framework
- Introduction to machine learning operations (MLOps)

- End-to-end machine learning operations (MLOps) with Azure Machine Learning

Skill Map for Azure AI Deployments

The mentee should deliberately connect AI concepts to these Azure service areas:

- Microsoft Foundry: models, agents, projects, evaluations, prompt flow
- Azure AI Search: grounding and RAG
- Azure Storage: document and data storage for prompts, content, logs, and artifacts
- Microsoft Entra ID: identity, RBAC, managed identities
- Azure Key Vault: secrets and key management
- Azure Monitor and Log Analytics: telemetry, alerting, KQL, operational visibility
- Azure Container Apps or App Service: host AI apps and APIs
- Azure Functions: event-driven orchestration where needed
- Azure Container Registry: container image management
- Virtual Network, Private Endpoints, and DNS` : secure enterprise connectivity
- GitHub Actions and Azure Machine Learning: automation and MLOps

Month 1: Build the AI Foundation on Top of Existing Azure Knowledge

Objective

Translate existing Azure infrastructure knowledge into AI solution vocabulary and understand where Microsoft Foundry fits in the Azure platform.

Official Microsoft Learn paths

- Introduction to AI in Azure
- Introduction to Cloud Infrastructure: Describe Azure architecture and services
- Introduction to Cloud Infrastructure: Describe Azure management and governance

Weekly plan

- Week 1: Learn AI concepts, machine learning basics, generative AI, and agents.
- Week 2: Refresh Azure compute, networking, storage, and identity with an AI workload mindset.
- Week 3: Study governance, compliance, cost management, and monitoring in Azure.
- Week 4: Write a 2-page summary answering: "What changes when infrastructure teams start hosting AI workloads?"

Expected outcomes

- Explain AI, ML, generative AI, agents, and RAG at architecture-review level.
- Map Azure core services to AI workload needs.
- Understand that AI workloads add governance, data, observability, and cost pressure beyond normal app hosting.

Hands-on focus

- Create a simple architecture sketch for an AI app using: Foundry, Storage, Entra ID, Key Vault, AI Search, and Azure Monitor.

Month 2: Get Productive with Microsoft Foundry

Objective

Learn how Microsoft Foundry is used to build AI apps and agents, and understand the common workload types available through the platform.

Official Microsoft Learn path

- Get started with AI applications and agents on Azure

Weekly plan

- Week 1: Work through the Foundry introduction and generative AI and agents modules.
- Week 2: Cover text analysis and speech modules.
- Week 3: Cover computer vision and information extraction modules.
- Week 4: Build a small demo backlog of 3 use cases the mentee could implement in Azure.

Expected outcomes

- Understand the functional breadth of Foundry: text, speech, vision, extraction, and agents.
- Identify which Azure AI capability matches which business problem.
- Gain confidence navigating the Foundry mental model before deeper engineering work.

Hands-on focus

- Prototype one simple Python-based use case: text summarization, document extraction, or speech-to-text.

Month 3: Build Real Generative AI Solutions with Microsoft Foundry

Objective

Move from exploration to implementation: model choice, SDK usage, prompt flow, RAG, evaluation, and responsible AI.

Official Microsoft Learn path

- Create custom copilots with Azure AI Foundry

Priority modules in this path

- Plan and prepare to develop AI solutions on Azure
- Choose and deploy models from the model catalog in Microsoft Foundry portal
- Develop an AI app with the Microsoft Foundry SDK
- Get started with prompt flow to develop language model apps in the Microsoft Foundry

- Develop a RAG-based solution with your own data using Microsoft Foundry
- Fine-tune a language model with Microsoft Foundry
- Implement a responsible generative AI solution in Microsoft Foundry
- Evaluate generative AI performance in Microsoft Foundry portal

Weekly plan

- Week 1: Study model selection, deployment patterns, and SDK-based development.
- Week 2: Learn prompt flow and prompt experimentation discipline.
- Week 3: Build a small RAG proof of concept with sample documents.
- Week 4: Run evaluation and document responsible AI controls and improvement actions.

Expected outcomes

- Understand the deployment lifecycle from model selection to evaluation.
- Know when to use prompting, RAG, or fine-tuning.
- Be able to explain the role of AI Search, storage, and data preparation in a Foundry solution.

Hands-on focus

- Deliverable: a small internal copilot that answers questions over a curated document set.

Month 4: Strengthen the Azure Infrastructure Layer for AI

Objective

Shift from app-building to platform-building: design AI-ready Azure infrastructure that is secure, observable, resilient, and cost-aware.

Official Microsoft Learn paths

- Manage AI-ready infrastructure
- Build great solutions with the Microsoft Azure Well-Architected Framework

Weekly plan

- Week 1: Study Foundry hub-and-project architecture, shared connections, and governance patterns.
- Week 2: Focus on monitoring, alerting, dashboards, and KQL-driven operations.
- Week 3: Cover resilience, security, cost optimization, and operational excellence using the Well-Architected Framework.
- Week 4: Redesign the Month 3 solution as an enterprise-ready target architecture.

Expected outcomes

- Understand how to structure hubs, projects, and shared services for team-scale AI work.
- Apply Azure Monitor, RBAC, identity, and policy concepts to AI systems.
- Be able to reason about cost, networking, and high availability for AI deployments.

Hands-on focus

- Produce an architecture document that includes: private access strategy, managed identity usage, Key Vault integration, monitoring, logging, and cost controls.

Month 5: Add Automation, DevOps, and MLOps Discipline

Objective

Treat AI workloads as production systems with repeatable deployment and lifecycle controls.

Official Microsoft Learn paths

- Introduction to machine learning operations (MLOps)
- End-to-end machine learning operations (MLOps) with Azure Machine Learning

Weekly plan

- Week 1: Learn the MLOps fundamentals: source control, automation, environments, and deployment thinking.
- Week 2: Study GitHub Actions for AI and ML workflows.
- Week 3: Review linting, unit testing, branch protection, and deployment automation.
- Week 4: Define a CI/CD flow for the Month 3 or Month 4 solution.

Expected outcomes

- Understand how AI systems move from experiment to managed release.
- Be able to describe where GitHub Actions, Azure Machine Learning jobs, and release environments fit.
- Build a repeatable promotion model for dev, test, and prod AI workloads.

Hands-on focus

- Create a deployment checklist covering: source control, secrets, environments, approvals, testing, rollback, and observability.

Month 6: Capstone Integration and Portfolio Outcome

Objective

Consolidate the learning into one end-to-end Microsoft Foundry on Azure deployment blueprint and demo solution.

Capstone scenario

Build a production-style enterprise knowledge assistant using:

- Microsoft Foundry for model and agent capabilities
- Azure AI Search for retrieval
- Azure Storage for documents
- Microsoft Entra ID for access control

- Azure Key Vault for secrets
- Azure Monitor and Log Analytics for telemetry
- App Service or Azure Container Apps for hosting
- GitHub Actions for CI/CD

Weekly plan

- Week 1: Finalize use case, scope, architecture, and backlog.
- Week 2: Implement the solution skeleton and deployment topology.
- Week 3: Add telemetry, security controls, and operational documentation.
- Week 4: Demo the solution and complete a formal retrospective.

Final deliverables

- Architecture diagram
- Service selection rationale
- Security and governance notes
- Monitoring and cost plan
- CI/CD flow description
- Demo script
- Lessons learned and next-step plan

Suggested Monthly Mentor Checkpoints

- End of Month 1: Can the mentee explain AI concepts in Azure infrastructure language?
- End of Month 2: Can the mentee identify which Foundry capability fits which workload?
- End of Month 3: Can the mentee build a basic generative AI or RAG prototype?
- End of Month 4: Can the mentee design AI-ready Azure infrastructure with governance and observability?
- End of Month 5: Can the mentee describe a production deployment flow and control model?
- End of Month 6: Can the mentee present an end-to-end Azure AI solution with Foundry and surrounding Azure services?

Recommended Mentor Emphasis Areas

- Push beyond theory: every month should end with a small artifact, not just completed modules.
- Keep the learner in architecture mode: always ask why a service is needed and what risk it addresses.
- Tie new AI learning back to existing strengths: networking, identity, monitoring, governance, and automation are major differentiators for infrastructure professionals moving into AI.
- Avoid overfocusing on model training early: for this learner, deployment architecture, governance, and operationalization are the highest-value first moves.

Optional Certification Direction

If the mentee wants exam-aligned outcomes after this roadmap:

- Azure fundamentals refresh: align with AZ-900 style infrastructure grounding
- AI foundation: align with the current Microsoft Learn AI introductory path
- Longer-term role growth: target Azure AI engineering and solution architecture capabilities after the capstone

Success Criteria at 6 Months

The roadmap is successful if the mentee can:

6. Design an Azure AI solution centered on Microsoft Foundry.
7. Explain the role of AI Search, identity, storage, monitoring, and deployment hosting around that solution.
8. Build or supervise a small RAG or agent-based proof of concept.
9. Define governance, security, cost, and observability requirements for production readiness.
10. Present a practical Azure-first AI deployment roadmap for a real business use case.