DevOpsSchool

Course curriculum · November 2026

LEVEL 04 OF 04 · ADVANCED

Azure Openai Advanced Training

Design resilient, governed, enterprise-scale Azure OpenAI platforms with an AI gateway, multi-region failover, hardened safety and LLMOps.

LIVE ONLINE **Max 8-10 hours**	CORPORATE (DAY-WISE) **5 days · Days 1-5**	SELF-PACED VIDEO **1-year LMS access**
BATCH SIZE **Max 10 learners**	NEXT BATCHES **1st of every month · weekends: first Saturday**	CERTIFICATE **Verifiable certificate of completion**

Foundation	Essential	Intermediate	Advanced
1 day · Included	2 days · Included	3 days · Included	5 days · This curriculum

INSIDE THIS CURRICULUM

What you will achieve	Outcomes and level focus
Who should attend	Audience and prerequisites
Agenda	50 topics · Foundation to Advanced
Hands-on labs	4 labs
Formats and fees	Live, 1:1, self-paced, corporate
Trainer and reviews	Rajesh Kumar
Enrolment	Batches, payment, contacts

WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

Course page: www.devopsschool.com/courses/openai/azure-openai-advance-training.html
Talk to us: WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

OUTCOMES

What you will achieve

By the end of the Azure Openai course path, you will be able to:

1

Use Azure OpenAI responsibly

Choose models, write effective prompts and apply content filtering and responsible AI policies.

2

Build generative AI apps

Develop chat, RAG, embedding, image, speech and agent features with Azure OpenAI APIs.

3

Run AI securely at scale

Secure, monitor, govern and scale Azure OpenAI with private networking and an AI gateway.

WHAT THE ADVANCED LEVEL ADDS

Resilience, AI gateway, governance and LLMOps

Hands-on: Capstone: enterprise AI assistant platform. Front two regional deployments with an API Management AI gateway, run a secured RAG agent behind it, enforce policy and evaluation gates, and test regional failover.

WHO SHOULD ATTEND

- Experienced Azure Openai practitioners responsible for quality, security and scale.
- Senior engineers, architects and leads who design and troubleshoot Azure Openai solutions.
- Teams making Azure Openai an organisation-wide standard.

BEFORE YOU START

- The Intermediate level, or equivalent hands-on experience.
- You can already: Build RAG and agent solutions, fine-tune models, and run Azure OpenAI securely with identity, private networking, monitoring, quota and cost control.
- Basic computer skills and a laptop with a stable internet connection for the live sessions and labs.
- This curriculum includes every earlier level, so it is complete on its own.

AGENDA · FOUNDATION TO ADVANCED

Azure Openai Advanced Training agenda

50 topics. Each level builds on the one before it, so this curriculum lists every earlier level in full. Days are for corporate (day-wise) delivery; live online batches cover the same agenda in max 8-10 hours.

DAY 1 — Foundation: Generative AI concepts, models, safety and prompting

Included from Foundation

1. What is generative AI: large language models, tokens, context windows and foundation models
2. What is Azure OpenAI: describing the service, its relation to OpenAI and how to use it via portal, SDKs and REST
3. Models: GPT-4o, GPT-4.1 and GPT-5 families, o-series reasoning, embeddings, image models and legacy GPT-3, GPT-35-Turbo, DALL-E
4. Capabilities: OpenAI natural language, code generation and image-generation capabilities with examples
5. Quotas and limits: tokens and requests per minute, regional quotas and model availability by region
6. Pricing: input and output tokens, standard versus provisioned throughput and batch processing
7. Quick start: create a resource, deploy a chat model and test it in the Azure AI Foundry playground
8. Access and responsible AI policies: limited access, code of conduct, transparency note, data, privacy and security
9. Abuse monitoring and content filtering: harm categories, severity levels and Prompt Shields
10. Embeddings concept: vectors, similarity and search, clustering and recommendation use cases
11. Prompt engineering: intro and techniques, zero-shot, few-shot, step-by-step reasoning, delimiters and output formats
12. System message design: role, scope, tone, safety instructions and grounding rules
13. Red teaming large language models: planning, harm categories and tools such as PyRIT
14. Knowledge check and summary: reviewing service concepts, safety controls and prompting skills

DAY 2 — Essential: Deployments, SDKs, chat, embeddings and filters

Included from Essential

15. Resource creation and model deployment: portal and Azure CLI, Global Standard, Data Zone, Standard and Provisioned types
16. Chat completions: GPT-35-Turbo and GPT-4 class models, messages and roles, temperature, top_p and max tokens
17. API styles: legacy Completions, Chat Completions and the Responses API with api-version choices
18. OpenAI vs Azure OpenAI in Python: AzureOpenAI client, endpoints, deployments, API keys and Entra ID tokens
19. Configure content filters: custom filter configurations, severity thresholds, blocklists and annotations
20. Work with code: code generation, explanation and refactoring with GPT models, successors to the retired Codex models
21. Function calling and structured outputs: tool definitions, JSON schema responses and parsing
22. Embeddings tutorial: text-embedding-3 models, chunking text and cosine similarity search
23. Speech to speech chat tutorial: Azure AI Speech with a GPT model and the gpt-4o realtime API
24. Image generation: DALL-E 3 and gpt-image-1 deployments, sizes, quality and edits
25. Robust calls: streaming responses, token counting with tiktoken and retries for 429 errors

DAY 3 — Intermediate: RAG, fine-tuning, agents, security and operations

Included from Intermediate

26. Retrieval-augmented generation: Azure AI Search vector and hybrid search, semantic ranker and citations
27. Use with large datasets: chunking strategies, indexing pipelines and the Batch API for bulk jobs

28. Prepare dataset and fine-tune your model: JSONL chat format, training and validation files, deploying the fine-tuned model
29. Manage identity: Entra ID authentication, managed identities and Cognitive Services OpenAI roles
30. Manage quota: TPM allocation per deployment, dynamic quota and provisioned throughput units
31. Use virtual networks: private endpoints, disabling public access and trusted service access
32. Encryption of data at rest: Microsoft-managed versus customer-managed keys in Key Vault
33. Monitor Azure OpenAI: Azure Monitor metrics, diagnostic logs and Application Insights tracing
34. Plan and manage costs: cost analysis, budgets, prompt caching and token-efficient prompts
35. Evaluations: groundedness, relevance, coherence and safety evaluators in Azure AI Foundry
36. Agents: Azure AI Foundry Agent Service with file search, code interpreter and function tools
37. App frameworks: Semantic Kernel or LangChain apps hosted on Azure Functions or App Service
38. Troubleshooting: 429 throttling, content filter blocks, context length errors and regional model availability

DAYS 4-5 **Advanced: Resilience, AI gateway, governance and LLMOps** **New at this level**

39. Business continuity and disaster recovery: multi-region deployments, Global and Data Zone types and failover design
40. AI gateway with API Management: load balancing across deployments, token limit policies, semantic caching and token metrics
41. Provisioned throughput at scale: PTU capacity planning, spillover to standard deployments and utilisation tracking
42. Enterprise landing zone: hub and spoke networking, private DNS, shared AI services and per-team access
43. Safety hardening: jailbreak and indirect prompt injection defences, groundedness detection and Defender for Cloud AI alerts
44. Governance: Azure Policy for AI services, approved model lists, data residency and audit logging
45. LLMOps: prompt versioning, evaluation gates in CI/CD, model version upgrades and retirement planning
46. Multi-agent systems: orchestration with Semantic Kernel or AutoGen and tool integration over MCP
47. Advanced customisation: distillation and reinforcement fine-tuning where available, comparing with RAG
48. Cost optimisation: model routing, smaller models for simple tasks, batch processing and caching
49. Responsible AI programme: impact assessments, transparency notes for your app, human oversight and EU AI Act duties
50. Capstone: an enterprise assistant with RAG, agents, AI gateway, private networking, evaluations and DR

PRACTICE

Hands-on labs

You build these yourself during the course, guided by the trainer.

FOUNDATION

Deploy a model and shape it

Create an Azure OpenAI resource, deploy a chat model, and iterate on a system message and few-shot prompt in the playground while observing content filter results.

ESSENTIAL

Build a grounded chat assistant

Write a Python app with the AzureOpenAI client using Entra ID, add function calling and embedding search over a small document set, and apply a custom content filter.

INTERMEDIATE

Secure RAG app on private network

Index documents in Azure AI Search, build a RAG endpoint with managed identity and private endpoints, evaluate groundedness and monitor token usage and cost.

ADVANCED

Capstone: enterprise AI assistant platform

Front two regional deployments with an API Management AI gateway, run a secured RAG agent behind it, enforce policy and evaluation gates, and test regional failover.

HOW YOU CAN LEARN

Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

FORMAT	DURATION	WHAT YOU GET	FEE
Live & interactive (online)	Max 8-10 hours	Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month.	**INR 34,999**
1:1 live training (online)	Max 8-10 hours, one-to-one	The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace.	**INR 99,999**
Self-paced video	1-year LMS access	All four Azure Openai levels, plus 20+ more courses and 50+ tools, at your own pace.	**INR 833/month** billed yearly at INR 9,996
Corporate (day-wise)	5 days for this level	A private batch for your team, on demand. Foundation 1 day, Essential 2, Intermediate 3, Advanced 5 (cumulative).	**Quote on request**

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on labs and assignments at every level. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

20 years in industry	**200+** corporate trainings
10,000+ engineers mentored	**70+** organisations coached

REVIEWS

What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR AZURE OPENAI TRAINING

Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All levels of this course: www.devopsschool.com/courses/openai/

ENROL OR RESERVE A SEAT
Azure Openai Advanced Training

WHATSAPP · INDIA
+91 99057 40781

WHATSAPP · USA
+1 (469) 756-6329

EMAIL
contact@DevOpsSchool.com