training code: AI-200 / ENG DL 5d / EN

Develop AI cloud solutions on Azure

This course teaches developers how to create, monitor, and troubleshoot AI solutions on Microsoft Azure. Students will learn how to implement Azure compute and containerization patterns to host applications, build serverless APIs with Azure Functions, and integrate services using event-driven and message-based architectures such as Azure Service Bus and Event Grid. The course also covers working with Azure data services that support AI workloads, including designing and querying solutions with Cosmos DB for NoSQL, Azure Database for PostgreSQL with pgvector, and Azure Managed Redis for caching, streaming, and vector search. By the end of the course, developers will be able to connect services, orchestrate AI workflows, and build secure, scalable, and observable AI-driven applications on Azure.

Training recipients

The training is intended for:

- Developers

Benefits

By completing this course, students will achieve the following objectives:

- Use Azure Container Registry to store and organize images, build images in the cloud with ACR Tasks, and apply tagging and versioning practices that support reliable deployments. From there, you deploy custom containers to Azure App Service, configure runtime behavior such as ports, startup commands, and persistent storage, and externalize environment-specific configuration using application settings.

- Start by deploying container apps to environments, configuring runtime settings with environment variables and secrets, and setting up registry authentication. You then manage the day-two lifecycle by updating images, managing revisions, monitoring logs, and configuring health probes. Finally, you learn to configure automatic horizontal scaling using HTTP rules, KEDA scalers, and traffic management to optimize performance and cost.
- Start by creating deployment manifests and exposing applications with Kubernetes Services. You then externalize configuration using ConfigMaps, secure sensitive settings with Secrets, and attach persistent storage for stateful workloads. Finally, you learn to monitor application health using logs and metrics, troubleshoot pod and Service issues, and verify connectivity paths to ensure reliable access to your applications.
- Start by building a data foundation with the Cosmos DB resource model, SDK integration, CRUD operations, and SQL queries to retrieve document data for AI applications.You then implement vector search capabilities to store embeddings, execute similarity queries using the VectorDistance function, combine vector search with metadata filters and hybrid search, and use the change feed to keep embeddings synchronized.Finally, you optimize query performance by analyzing query patterns, configuring range and composite indexes, selecting vector index types, and choosing consistency levels that balance freshness with cost efficiency.
- Start by building a data foundation with schema design, efficient SQL queries, and secure Python integration using Microsoft Entra authentication.You then implement vector search using the pgvector extension to store embeddings, execute similarity searches with different distance metrics, and build retrieval patterns that integrate with RAG pipelines for semantic search and recommendations.Finally, you optimize vector search performance by tuning PostgreSQL and pgvector configuration, selecting appropriate vector indexes, designing efficient data layouts, scaling for high-volume workloads, and implementing connection pooling for AI applications.

Training program

1.Implement container application hosting on Azure
- Store and manage containers in Azure Container Registry
- Deploy containers to Azure App Service

2.Deploy and manage apps on Azure Container Apps
- Deploy containers to Azure Container Apps
- Manage containers in Azure Container Apps
- Scale containers in Azure Container Apps

3.Deploy and monitor applications on Azure Kubernetes Service
- Deploy applications to Azure Kubernetes Service
- Configure applications on Azure Kubernetes Service
- Monitor and troubleshoot applications on Azure Kubernetes Service

4.Develop AI solutions with Azure Cosmos DB for NoSQL

- Build queries for Azure Cosmos DB for NoSQL
- Implement vector search on Azure Cosmos DB for NoSQL
- Optimize query performance for Azure Cosmos DB for NoSQL

5.Develop AI solutions with Azure Database for PostgreSQL

- Build and query with Azure Database for PostgreSQL
- Implement vector search with Azure Database for PostgreSQL
- Optimize vector search in Azure Database for PostgreSQL

Expected preparation of the participant

- Basic understanding of Azure services and cloud computing concepts.
- Familiarity with container concepts.
- Familiarity with command-line tools including Azure CLI.
- Familiarity with containerization and Kubernetes fundamentals.
- Programming experience with Python.
- Familiarity with JSON document structures.
- Familiarity with relational databases and SQL fundamentals.
- Understanding of machine learning concepts including embeddings and similarity search.

Training Includes

- manual in electronic form available on the platform: https://learn.microsoft.com/pl-pl/training/
- access to Altkom Akademia's student portal

Duration

5 days / 35 hours

Language

- Training: English
- Materials: English