

AI in Microsoft SQL Server

Course code: GOC633

This course is designed for database developers and architects who want to make effective use of artificial intelligence directly within Microsoft SQL Server 2025 and Azure SQL Database. You will learn to work with machine learning models built in Python and integrate them into your databases' stored procedures — we will show you how to train your own predictive model directly on the data in your database and how to expose its predictions to applications as an ordinary stored procedure. You will get familiar with the new vector data type and DiskANN vector indexes and discover how they let you connect large language models to your database and search by the meaning of text rather than by mere keyword matching. You will learn to combine full-text and vector search into hybrid search and to build a complete RAG solution ("ask your data") purely in T-SQL — both with the Azure OpenAI service and with a locally running model, without your data ever leaving the corporate network. You will find out how GitHub Copilot in SQL Server Management Studio and Visual Studio Code helps you write SQL queries and develop databases, and how to align its behavior with your team standards. Because AI-generated code needs firm guardrails, you will also learn to manage your database as code using SQL Server Database Projects — from schema versioning in Git through pull requests and automated builds to CI/CD deployment, schema drift detection, and database testing. Finally, we will show you how to turn SQL Server into an MCP server that gives AI agents secure and auditable access to your data.

Affiliate	Duration	Course price	ITB
Praha	3	21 000 Kč	30
Brno	3	21 000 Kč	30
Bratislava	3	875 €	30

The prices are without VAT.

Course terms

Date	Duration	Course price	Type	Course language	Location
05.10.2026	3	21 000 Kč	Presence	CZ/SK	GOPAS Praha
19.10.2026	3	875 €	Presence	CZ/SK	GOPAS Bratislava
18.11.2026	3	21 000 Kč	Presence	CZ/SK	GOPAS Brno
21.12.2026	3	21 000 Kč	Online	CZ/SK	Online
21.12.2026	3	875 €	Online	CZ/SK	Online

The prices are without VAT.

What we will teach you

- You will get familiar with the complete range of AI features in Microsoft SQL Server 2025 and find out which of them are ready for production deployment
- You will learn to integrate machine learning models built in Python into stored procedures and train your own predictive model directly on the data in your database
- You will try out semantic search by the meaning of text using the vector data type, embeddings, and DiskANN vector indexes
- You will learn to combine full-text and vector search into hybrid search
- You will try out building a complete "ask your data" RAG solution purely in T-SQL — using Azure OpenAI as well as a locally running model without sending data outside the corporate network
- You will learn to use GitHub Copilot effectively when writing SQL queries and developing databases, and to align its behavior with team standards
- You will learn to manage a database as code with SQL Server Database Projects, including automated CI/CD deployment and testing
- You will find out how to turn SQL Server into an MCP server and expose data to AI agents securely

GOPAS Praha
Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno
Nové sady 996/25
602 00 Brno
Tel.: +420 530 513 590
info@gopas.cz

GOPAS Bratislava
Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk


Copyright © 2026 GOPAS, a.s.,
All rights reserved

AI in Microsoft SQL Server

Who the course is intended for

- Database developers who want to enrich their applications with predictions, semantic search, and answers generated by language models over company data.
- Developers building applications on Microsoft SQL Server or Azure SQL Database who want to speed up development with AI assistants and adopt modern database DevOps practices.
- Architects and technically oriented database administrators who design secure ways of bringing artificial intelligence and AI agents into solutions built on company data.

Prerequisites

- Knowledge of T-SQL at the level of the GOC278 course.
- Knowledge of database development in Microsoft SQL Server at the level of the GOC632 course.
- Basic knowledge of Python is an advantage (not required; the exercise code will be provided ready to use).
- Basic knowledge of working with Git is an advantage.

Course outline

1 Overview of AI Capabilities in Microsoft SQL Server

- You will get familiar with the complete range of AI features in SQL Server 2025 and find out which computations run directly inside the database engine and which rely on external services
- You will understand the differences in AI capabilities between SQL Server 2025, Azure SQL Database, and SQL database in Microsoft Fabric
- You will find out which features are already generally available and which are still in preview, and what that means for production deployment
- You will learn to assess which data may leave the database or the corporate network, and why data classification is the first step of every AI solution
- You will discover when integrating AI directly into the database makes sense and when a different architecture is the better choice

2 Machine Learning in SQL Server – Python in Stored Procedures

- You will understand the architecture of Machine Learning Services and how SQL Server executes external scripts
- You will learn to install and configure a Python runtime on SQL Server 2022/2025 and avoid the common installation pitfalls
- You will learn to run Python code from T-SQL using the `sp_execute_external_script` procedure and pass data between SQL Server and Python
- You will try out training your own machine learning model directly on the data in your database and storing it in a table
- You will learn to expose the model's predictions to applications as an ordinary stored procedure with no changes to application code
- You will find out when to use native scoring with the `PREDICT` function and how it differs from external scripts in both performance and limitations
- You will discover how to govern the resources allocated to external scripts and when it is better to train models outside the database

3 Vectors and Embeddings – Fundamentals

- You will understand what tokens and embeddings are and how they make it possible to compare texts by their meaning
- You will get familiar with the new vector data type, its dimensions, and its storage requirements
- You will learn to manage credentials for AI services securely — from Managed Identity in Azure to API keys in on-premises environments
- You will learn to register external AI models with the `CREATE EXTERNAL MODEL` statement and connect SQL Server both to Azure OpenAI and to a locally running model, without data leaving the corporate network

GOPAS Praha
Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno
Nové sady 996/25
602 00 Brno
Tel.: +420 530 513 590
info@gopas.cz

GOPAS Bratislava
Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk



Copyright © 2026 GOPAS, a.s.,
All rights reserved

AI in Microsoft SQL Server

- You will try out generating embeddings with the AI_GENERATE_EMBEDDINGS function and splitting long texts into chunks with the AI_GENERATE_CHUNKS function
- You will find out which data is worth converting into embeddings and how the choice of chunk size affects search quality

4 Vector Search and Scaling It

- You will learn to search by similarity with the VECTOR_DISTANCE function and choose an appropriate vector comparison metric
- You will understand how the DiskANN vector index and approximate search with the VECTOR_SEARCH function work
- You will try out creating a vector index and measuring the difference between exact and approximate search
- You will find out what limitations vector indexes have, how they are evolving, and what can be taken to production today
- You will discover how to keep embeddings up to date as source data changes — from triggers through Change Tracking and Change Data Capture to event streaming

5 Intelligent Search – Full-Text, Vectors, and Hybrid

- You will learn the fundamentals of full-text search: the full-text index and the CONTAINS, FREETEXT, and CONTAINSTABLE predicates
- You will understand the strengths and weaknesses of full-text and vector search and learn to choose the right approach for a given scenario
- You will learn to combine both approaches into hybrid search using the Reciprocal Rank Fusion algorithm
- You will try out comparing all three search methods on the same data and evaluating the quality of their results

6 RAG Directly in T-SQL

- You will understand the Retrieval Augmented Generation (RAG) pattern and find out when to use it instead of fine-tuning a model
- You will learn to prepare context for a language model using FOR JSON and decide which data to send to the model and which to hold back
- You will learn to build a prompt directly in T-SQL with the JSON_OBJECT and JSON_ARRAY functions, including rules for grounding answers in your data
- You will learn to call language models with the sp_invoke_external_rest_endpoint procedure, including error handling and retrying failed calls
- You will try out building a complete "ask your data" RAG solution in a single stored procedure — against both Azure OpenAI and a local model
- You will find out how to secure the RAG procedure, control access and costs with permissions, and monitor AI service calls (latency, token consumption, error rates)

7 GitHub Copilot for Database Developers

- You will learn to use GitHub Copilot in SQL Server Management Studio and Visual Studio Code, including conversations over a connected database
- You will try out real-world workflows: generating procedures and views, explaining unfamiliar code, suggesting optimizations, and generating test data or documentation
- You will learn to align Copilot's behavior with your team standards using instruction files and prompt files
- You will find out what data Copilot sends to the cloud and learn to set it up for safe use, including review of the generated code
- You will discover where the limits of AI assistants lie and when you need to reach for traditional performance tuning methods

8 Database as Code – SQL Database Projects and DevOps for AI-Assisted Development

GOPAS Praha
Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno
Nové sady 996/25
602 00 Brno
Tel.: +420 530 513 590
info@gopas.cz

GOPAS Bratislava
Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk



Copyright © 2026 GOPAS, a.s.,
All rights reserved

AI in Microsoft SQL Server

- You will understand why AI-assisted development requires keeping the database in version control with automated validation of changes
- You will learn to create a modern SDK-style SQL Database Project, build it into a .dacpac package, and deploy it with the SqlPackage tool
- You will find out how to organize a database project in Git, including pre/post-deployment scripts and idempotent management of reference data
- You will learn a team workflow with feature branches and pull requests as the place to review AI-generated code, backed by automated build validation
- You will discover how to detect and resolve schema drift — differences between the project and the live database
- You will try out building a CI/CD pipeline that automatically builds and deploys schema changes to a database
- You will find out how to test a database, from build validation through unit tests to integration tests, and how Copilot can help you create the tests

9 SQL Server as an MCP Server – Secure AI Agent Access to Data

- You will understand the principles of the Model Context Protocol (MCP) and find out why AI agents need governed tools instead of freely generated SQL queries
- You will learn to build a SQL MCP Server on top of the Data API builder tool and expose selected database entities to AI agents
- You will try out visual configuration in the MSSQL extension for Visual Studio Code and see how a single configuration can expose REST, GraphQL, and MCP endpoints at the same time
- You will learn to secure agentic access to data: entity-level permissions, read-only mode, excluded tables, defense against prompt injection, and auditing
- You will try out connecting an AI agent to the database and verifying that the security guardrails you configured actually hold
- You will find out where to go next — Fabric data agents, Microsoft Foundry, and building your own AI agents

GOPAS Praha
Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno
Nové sady 996/25
602 00 Brno
Tel.: +420 530 513 590
info@gopas.cz

GOPAS Bratislava
Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk



Copyright © 2026 GOPAS, a.s.,
All rights reserved