



Document Generated: 07/03/2026

Learning Style: Virtual Classroom

Technology: Microsoft

Difficulty: Intermediate

Course Duration: 4 Days

Develop AI-enabled Database Solutions (DP-800)



About This Course:

It is intended for professionals who build modern data solutions that integrate structured and semi structured data and incorporate AI features into scalable enterprise applications. It will also be valuable for individuals who develop applications that rely on SQL based data services enhanced with vector search, embeddings, and other AI driven capabilities.

Course Objectives:

- Design and implement database objects with SQL
- Implement programmability objects with SQL
- Write advanced T-SQL code
- Implement SQL solutions by using AI-assisted tools
- Implement data security and compliance with SQL
- Optimize database performance
- Implement CI/CD by using SQL Database Projects
- Integrate SQL solutions with Azure services
- Design and implement models and embeddings with SQL
- Design and implement intelligent search with SQL
- Design and implement RAG with SQL

Audience:

- The audience for this course is data professionals who want to learn about designing and developing AI-enabled database solutions across Microsoft's SQL platforms, including SQL Server, Azure SQL, and SQL databases in Microsoft Fabric. This role develops database solutions that include both structured and semi-structured data and integrates AI features into modern and highly scalable enterprise applications.

Prerequisites:

- None

Course Outline:

- Understand your SQL Server-based platform choices
- Build effective tables
- Optimize with indexes

- Use specialized table types
- Enforce data integrity with constraints
- Manage JSON columns and indexes
- Partition tables for scale
- Module assessment
- Create views
- Create stored procedures
- Create scalar functions
- Create table-valued functions
- Create triggers
- Choose when to use each option
- Organize queries with Common Table Expressions
- Apply window functions for analytics
- Process JSON data with built-in functions
- Match patterns with regular expressions
- Find approximate matches with fuzzy string functions
- Traverse relationships with graph queries
- Compare rows with correlated subqueries
- Handle errors with TRY...CATCH
- Module assessment
- Describe AI-assisted development tools available for Microsoft SQL platforms
- Interpret security impact of using AI-assisted tools
- Enable GitHub Copilot and Fabric Copilot
- Configure model and Model Context Protocol (MCP) tool options in a GitHub Copilot or Fabric Copilot chat session
- Create and configure GitHub Copilot instruction files
- Connect to MCP server endpoints, including Microsoft SQL Server and Fabric Lakehouse
- Module assessment
- Protect data with encryption
- Configure dynamic data masking
- Implement row-level security
- Manage permissions and secure access
- Implement auditing
- Configure secure access to AI services
- Secure data API endpoints
- Module assessment
- Recommend database configurations
- Preserve data integrity with transaction isolation levels and concurrency controls
- Evaluate query performance with execution plans and DMVs
- Monitor and tune queries with Query Store and Query Performance Insight
- Identify and resolve blocking and deadlocks
- Create, build, and validate SQL Database Projects
- Configure source control and manage reference data
- Manage branching, pull requests, and conflict resolution
- Detect and resolve schema drift
- Implement CI/CD pipelines
- Design and implement a testing strategy

- Create configuration files for Data API Builder
- Define entities for REST and GraphQL
- Expose database objects, stored procedures, and views
- Explore deployment options for Data API Builder
- Recommend Azure Monitor configurations
- Handle changes with event-driven patterns
- Module assessment
- Understand and evaluate models for SQL database workloads
- Create and manage external models in SQL
- Design embeddings for SQL database workloads
- Generate and maintain embeddings for SQL database workloads
- Choose an intelligent search approach
- Implement full-text search
- Prepare SQL for vector search
- Implement vector search query patterns
- Implement hybrid search and ranking
- Identify RAG use cases and architecture
- Prepare retrieval context for augmentation
- Augment prompts with database context
- Generate and process RAG responses