

Data Engineering v Microsoft Fabric

Course code: MOC DP-700

This course is designed for data professionals experienced in data extraction, transformation and loading who want to implement complete data engineering solutions on the Microsoft Fabric platform. You will learn to design and develop data loading processes, data architectures and their orchestration. You will practice ingesting and transforming data using Dataflows Gen2, pipelines and Apache Spark notebooks, storing data in a lakehouse as Delta Lake tables and organizing it according to the medallion architecture. You will learn to work with a data warehouse, load data into it and query it using T-SQL. You will also be introduced to real-time data processing with eventstreams, an eventhouse and the KQL language, to building real-time dashboards and to automating actions with Activator. The final part of the course covers deploying changes through CI/CD, monitoring, securing and administering the entire Microsoft Fabric environment. The course also serves as preparation for the DP-700 certification exam.

What you will learn

- You will get to know the Microsoft Fabric platform and its data engineering components
- You will learn to ingest and transform data using Dataflows Gen2, pipelines and Apache Spark
- You will practice storing data in lakehouse and data warehouse and working with Delta Lake tables
- You will learn to organize data in a lakehouse using the medallion architecture
- You will find out how to process real-time data using eventstreams, an eventhouse and the KQL language
- You will practice automating the deployment of changes with CI/CD, Git integration and deployment pipelines
- You will learn to monitor and secure an analytics solution in Microsoft Fabric

Who the course is intended for

- Data engineers who want to design and develop data ingestion and transformation processes, data architectures and their orchestration in the Microsoft Fabric environment.
- Data professionals experienced in data extraction, transformation and loading who want to extend their skills to the Microsoft Fabric platform.
- Anyone preparing for the DP-700: Implementing Data Engineering Solutions Using Microsoft Fabric certification exam.

Prerequisites

- Experience with data extraction, transformation and loading (ETL/ELT processes)
- Experience manipulating and transforming data in at least one of the following languages: SQL, PySpark or KQL
- Basic knowledge of core data concepts and Microsoft Azure data services at the level of the MOC DP-900 course is recommended

Course outline

1 End-to-end analytics with Microsoft Fabric

- You will get to know the Microsoft Fabric platform, OneLake storage and the individual workloads: Data Factory, Data Engineering, Data Warehouse, Data Science, Real-Time Intelligence and Power BI
- You will find out which data team roles work with which services in Microsoft Fabric
- You will discover how to enable Microsoft Fabric in your tenant and how to assign workspaces to a compute capacity

2 Get started with lakehouses in Microsoft Fabric

- You will understand what a lakehouse is and how it combines the flexible storage of a data lake with the analytical capabilities of a data warehouse
- You will learn to create your own lakehouse, including the automatically provisioned SQL analytics endpoint and default semantic model
- You will practice ingesting data into files and tables in a lakehouse and using shortcuts to access external data without copying it
- You will learn to query lakehouse tables using SQL

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

Data Engineering v Microsoft Fabric

- You will get to know the tools for exploring, transforming and visualizing data in a lakehouse

3 Use Apache Spark in Microsoft Fabric

- You will get to know the Apache Spark framework and the principles of distributed data processing in Spark pools
- You will learn to configure the Spark environment, including runtime versions, libraries and properties
- You will practice running Spark code interactively in a notebook and non-interactively using a Spark job definition
- You will learn to load, transform and save data in Spark dataframes and query it using Spark SQL
- You will discover how to visualize data directly in a notebook

4 Work with Delta Lake tables in Microsoft Fabric

- You will understand the principles of the Delta Lake format and its role in Microsoft Fabric
- You will learn to save dataframes as managed and external delta tables and to modify the data in them
- You will practice processing streaming data with Spark Structured Streaming, using a delta table as both a source and a sink
- You will find out how to optimize delta tables using the OptimizeWrite and V-Order functions, the Optimize and Vacuum commands and file partitioning
- You will discover how data versioning and querying historical table versions (time travel) work

5 Ingest data with Dataflow Gen2

- You will understand the principles of Dataflow Gen2, along with its benefits and limitations
- You will learn to create a dataflow to extract and transform data in the graphical Power Query environment
- You will find out how to configure a data destination and how to integrate Dataflow Gen2 into a pipeline

6 Orchestrate processes and data movement

- You will get to know the capabilities of pipelines in Microsoft Fabric, their activities and parameterization
- You will learn to use the Copy Data activity to move data
- You will practice creating a pipeline based on predefined templates
- You will learn to validate, run, schedule and monitor pipeline runs

7 Organize a lakehouse using the medallion architecture

- You will understand the principles of the medallion architecture and the purpose of the bronze, silver and gold layers
- You will learn to design and implement the individual layers and to choose the right tools for moving and transforming data between them
- You will find out how to analyze data stored in a lakehouse using Direct Lake mode in Power BI
- You will discover how to secure the individual layers using workspace and item permissions

8 Get started with Real-Time Intelligence in Microsoft Fabric

- You will understand the concepts of real-time data analytics
- You will get to know the components of Real-Time Intelligence: eventstream, eventhouse, real-time dashboard and Activator
- You will find out how to discover and connect to real-time data sources using the Real-Time hub

9 Use eventstreams in Microsoft Fabric

- You will get to know the components of eventstreams and the supported sources and destinations for streaming data
- You will learn to configure eventstream sources and destinations
- You will practice capturing, transforming and routing real-time data using transformations such as filtering and aggregation

10 Work with real-time data in an eventhouse

- You will learn to create an eventhouse and a KQL database and to load both static and streaming data into tables
- You will practice querying real-time data with the Kusto Query Language (KQL)

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

Data Engineering v Microsoft Fabric

- You will learn to create materialized views and stored functions in a KQL database

11 Create real-time dashboards

- You will learn to create a real-time dashboard over a streaming data source
- You will find out how to organize and filter dashboard data using base queries, pages and parameters
- You will discover how to manage and optimize a dashboard according to best practices

12 Use Activator

- You will understand how Activator works with events, objects, properties and rules
- You will learn to define objects and their properties over data from an eventstream
- You will learn to create rules that evaluate conditions in your data
- You will practice configuring actions triggered when conditions are met, such as sending an email, posting a Teams message or running a Power Automate flow

13 Get started with data warehouses in Microsoft Fabric

- You will get to know the fundamentals of data warehouses in Microsoft Fabric and how they differ from a lakehouse
- You will learn to design a data warehouse schema, including fact tables, dimension tables and star and snowflake schemas
- You will practice creating a data warehouse, cloning tables and querying data using T-SQL and the visual query editor
- You will find out how to define relationships between tables, how the default semantic model works and how to visualize data
- You will discover how to control access to a data warehouse using workspace and item permissions

14 Load data into a data warehouse

- You will understand data loading strategies, including full and incremental loads and the use of staging tables
- You will learn to work with dimension keys and to implement slowly changing dimensions
- You will practice loading data using a pipeline with the Copy Data activity, T-SQL statements (COPY INTO, CREATE TABLE AS SELECT) and Dataflows Gen2
- You will find out how Copilot can assist you when transforming data

15 Monitor a data warehouse

- You will learn to monitor capacity unit consumption using the Microsoft Fabric Capacity Metrics app
- You will practice monitoring currently running queries using dynamic management views
- You will discover how to analyze querying trends using query insights views

16 Secure a data warehouse

- You will get to know the concepts of securing a data warehouse in Microsoft Fabric
- You will learn to implement dynamic data masking
- You will practice configuring row-level security and column-level security
- You will learn to configure granular permissions on data warehouse objects using T-SQL

17 Implement CI/CD in Microsoft Fabric

- You will understand the principles of continuous integration and continuous delivery and how they are implemented in Microsoft Fabric
- You will learn to connect a workspace to Azure DevOps or GitHub version control, commit and sync changes and work with branches
- You will practice automating the deployment of content between environments using deployment pipelines
- You will find out how to automate CI/CD processes using the Fabric REST APIs

18 Monitor activities in Microsoft Fabric

- You will get to know monitoring concepts and best practices in Microsoft Fabric
- You will learn to use the Monitoring hub to track the status and progress of activities across Fabric items

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

Data Engineering v Microsoft Fabric

- You will discover how to trigger automated actions on your data using Activator

19 Secure data access in Microsoft Fabric

- You will understand the Microsoft Fabric security model and the sequential evaluation of access starting with Microsoft Entra ID authentication
- You will learn to configure workspace roles and item-level permissions
- You will find out how to control read access to lakehouse data through the SQL endpoint and Apache Spark
- You will practice configuring granular permissions with T-SQL and securing folders and files using OneLake data access roles

20 Administer Microsoft Fabric

- You will get to know the architecture and concepts of Microsoft Fabric, including tenants, capacities and workspaces
- You will find out which tasks a Fabric administrator handles and which tools are available: the admin portal, the admin monitoring workspace and the Capacity Metrics app
- You will discover how to assign and manage user licenses
- You will find out about data governance options such as endorsing content, tracking data lineage and running metadata scanning

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