

Designing and Implementing a Modern Data Warehouse and ETL/ELT Processes in Microsoft Fabric

Course code: GOC683

This course teaches participants how to design and implement modern data warehouses in Microsoft Fabric and how to build robust ETL/ELT processes on top of them using Data Factory. The course combines proven dimensional modelling theory (star schema, dimension and fact table types, slowly changing dimensions) with its practical implementation in Fabric Data Warehouse, with a strong emphasis on T-SQL, which is the primary language for working with data in Fabric Warehouse. You will learn how to design a layered warehouse architecture, implement its schema in T-SQL and secure it at the level of the workspace, database objects and individual rows and columns. Considerable attention is paid to advanced Data Factory scenarios – parameterised and metadata-driven pipelines, error handling, run auditing and, in particular, design patterns for incremental data loading from transactional systems. The course concludes with operations, performance tuning, deploying changes and migrating existing solutions from Microsoft SQL Server and SQL Server Integration Services, whose concepts are continuously compared with their counterparts in Microsoft Fabric. The course deliberately focuses on technologies built around T-SQL; in selected scenarios, however, it shows where a short Python notebook is the simplest way to reach the goal (Python and PySpark are covered systematically in courses GOC681 and GOC685). Given the breadth of the topics covered, the course is taught at a brisk pace and does not repeat general Microsoft Fabric platform topics – entry-level knowledge equivalent to course GOC680 is therefore strongly recommended.

Affiliate	Duration	Course price	ITB
-----------	----------	--------------	-----

The prices are without VAT.

Course terms

Date	Duration	Course price	Type	Course language	Location
------	----------	--------------	------	-----------------	----------

The prices are without VAT.

What You Will Learn

- You will get familiar with the architecture of modern data warehouses and their implementation in Microsoft Fabric
- You will learn how to design dimensional models – fact tables, dimensions including slowly changing dimensions (SCD) and the choice of data granularity
- You will try out implementing a complete data warehouse schema in Fabric Warehouse using T-SQL
- You will learn how to build parameterised and metadata-driven pipelines in Data Factory, including error handling and auditing
- You will master data transformations in T-SQL, including loading SCD Type 1 and SCD Type 2 dimensions and fact tables
- You will learn how to design and implement incremental data loading from transactional systems
- You will find out how to secure and monitor a data warehouse and tune its performance
- You will discover how to deploy solution changes across environments and how to migrate existing solutions from Microsoft SQL Server and SSIS

Who Should Attend

- Data warehouse developers and architects who want to build solutions in Microsoft Fabric.
- Data warehouse specialists working with Microsoft SQL Server and ETL developers working with SQL Server Integration Services who want to transfer their know-how to Microsoft Fabric.
- BI developers and data engineers who prefer working in T-SQL.
- Anyone looking for a comprehensive, practically oriented introduction to data warehouse design and implementation.

Prerequisites

- Knowledge of the Microsoft Fabric platform equivalent to course GOC680 is strongly recommended – the course does not repeat general platform topics (capacities, licensing, working with Lakehouse and Apache Spark), and

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

Designing and Implementing a Modern Data Warehouse and ETL/ELT Processes in Microsoft Fabric

without this prior knowledge the pace of the course will be very demanding.

- Knowledge of T-SQL at the level of writing queries, ideally equivalent to course GOC278; experience with writing stored procedures is an advantage.
- Experience with data warehouses or SQL Server Integration Services is an advantage, but not a requirement.

Course Outline

1 The Modern Data Warehouse and an Introduction to Fabric Data Warehouse

- You will understand the role of the data warehouse in a modern data platform and its place among the Microsoft Fabric components
- You will get familiar with classic data warehouse architectures and the Kimball and Inmon approaches
- You will find out how to map the medallion architecture (bronze/silver/gold) onto Warehouse schemas or a combination of Lakehouse and Warehouse
- You will understand the shift from ETL to ELT and its impact on process design compared with the SSIS mindset
- You will get familiar with the key features of Fabric Data Warehouse and learn how to choose between Lakehouse and Warehouse
- You will try out the options for working with a Warehouse, from the web editor to SQL Server Management Studio
- You will discover how Fabric Warehouse differs conceptually from a data warehouse on Microsoft SQL Server and Azure Synapse

2 Dimensional Modelling of a Data Warehouse

- You will learn the dimensional modelling process – identifying business processes, choosing the grain and building a bus matrix
- You will understand the differences between star and snowflake schemas and when to use each of them
- You will learn how to design dimensions, including surrogate keys, the date dimension and slowly changing dimensions (SCD Type 0 to 3)
- You will get familiar with special dimension types – role-playing, junk, degenerate and mini-dimensions
- You will learn how to design fact tables of various types (transaction, snapshot, factless) and understand measure additivity
- You will find out how to handle unknown dimension members and late arriving data at the design level
- You will try out designing a complete dimensional model for a given business domain

3 Implementing the Data Warehouse Schema in Fabric Warehouse

- You will find out how T-SQL in Fabric Warehouse differs from Microsoft SQL Server – data types, collation, DDL limitations
- You will learn how to organise warehouse layers using database schemas and how to query data across databases
- You will understand the role of NOT ENFORCED constraints and master strategies for generating surrogate keys
- You will try out zero-copy table cloning and time travel queries for development, testing and auditing
- You will find out how statistics work and how data is physically stored in the Delta-parquet format
- You will try out implementing the schema of the designed dimensional model in T-SQL, including populating the date dimension

4 Introduction to Data Factory in Microsoft Fabric

- You will get familiar with the Data Factory components and their roles – Pipelines, Dataflows Gen2 and Copy Job
- You will understand the concept of pipeline activities and learn how to manage connections to data sources
- You will discover the strengths and weaknesses of Dataflows Gen2 compared with T-SQL transformations and when to reach for which tool
- You will find out how to connect to on-premises data sources through a data gateway
- You will learn how to schedule execution and monitor runs in the Monitor hub
- You will understand how SSIS concepts map onto their counterparts in Data Factory

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

Designing and Implementing a Modern Data Warehouse and ETL/ELT Processes in Microsoft Fabric

5 Advanced Work with Pipelines

- You will learn how to parameterise pipelines and use the expression language for dynamic content
- You will master control flow activities (ForEach, If Condition, Switch, Until) and the Lookup, Get Metadata and Script activities
- You will understand the metadata-driven data loading pattern with a control table of sources
- You will learn how to handle errors, configure retry policies and send notifications through Teams and Outlook
- You will try out building a metadata-driven pipeline for populating the staging layer, including error handling

6 Data Loading Techniques for the Data Warehouse

- You will explore the Copy activity in depth – connectors, column mapping and performance settings
- You will learn how to load files from OneLake in T-SQL using the COPY INTO statement and the OPENROWSET function
- You will find out how to load data with cross-database queries against the SQL analytics endpoint of a Lakehouse
- You will get familiar with Copy Job and database Mirroring as a near real-time data source
- You will discover in which scenarios a Python notebook is the simplest way forward – calling REST APIs or processing complex JSON data
- You will try out loading data from files in OneLake into staging tables and learn how to choose the right loading technique

7 Data Transformations in T-SQL and Loading Dimensions and Facts

- You will learn how to implement ELT transformations using stored procedures, CTAS, INSERT ... SELECT and the MERGE statement
- You will understand transactions and the isolation level in Fabric Warehouse and learn how to write idempotent procedures
- You will try out implementing the loading of SCD Type 1 and SCD Type 2 dimensions, including change detection
- You will learn how to load fact tables, including surrogate key lookups and handling unmatched dimension members
- You will find out how to cleanse and validate data in T-SQL and how to build a load auditing framework

8 Incremental Data Loading

- You will understand why full loads are not enough and what the patterns for incremental processing are
- You will learn how to implement watermark-based loading with a control table and dynamic queries
- You will find out how to leverage Change Tracking and Change Data Capture on SQL Server sources, including capturing deleted records
- You will discover how to copy data incrementally using Copy Job and how to use a mirrored database as a source of changes
- You will learn how to implement upserts at the destination, handle late arriving data and ensure repeatable runs
- You will try out implementing complete incremental loading of a transactional table

9 Orchestration, Monitoring and Performance Tuning

- You will learn how to design master orchestration with dependencies between dimensions and facts, parallelisation and restartability
- You will find out how to trigger processing both on a schedule and based on events, such as a file arriving in OneLake
- You will learn how to monitor the solution using the Monitor hub, Query Insights views and Dynamic Management Views
- You will discover how to tune data warehouse performance – statistics, result set caching, V-Order and data compaction
- You will find out how to track and optimise the capacity consumption of ETL processes
- You will try out building a master pipeline for a complete daily warehouse load and analysing its run

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

Designing and Implementing a Modern Data Warehouse and ETL/ELT Processes in Microsoft Fabric

10 Securing the Data Warehouse and Providing Data to Consumers

- You will understand the Fabric Warehouse security model – workspace roles, item sharing and how they combine with T-SQL permissions
- You will learn how to configure granular permissions in T-SQL at the database, schema, object and column level
- You will try out implementing Row-Level Security and Dynamic Data Masking to protect sensitive data
- You will find out how to provide data to consumers through the SQL analytics endpoint and external tools
- You will get familiar with recommendations for designing the gold layer as a source for semantic models and reporting

11 Deploying Solutions and Migrating from Microsoft SQL Server and SSIS

- You will learn how to manage the solution lifecycle using Git integration and deploy changes through deployment pipelines
- You will find out how to work with a Warehouse as a database project and manage schema changes
- You will discover how to assess and carry out the migration of a data warehouse from Microsoft SQL Server, including the use of the Migration Assistant
- You will understand how to map SSIS processes onto Fabric components and choose a phased migration strategy
- You will try out connecting a workspace to Git and deploying a schema change to another environment
- You will get familiar with best practices and the most common mistakes when building a data warehouse in Fabric

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