



Microsoft Fabric Training: Lakehouse, OneLake and Analytics Engineering

27 September – 1 October 2027
London - Central London four-star



Microsoft Fabric Training: Lakehouse, OneLake and Analytics Engineering

Ref: 307_12361 **Date:** 27 September – 1 October 2027 **Location:** London - Central London four-star **Fees:** 23000 SAR

Introduction:

Many organisations adopt Microsoft Fabric but keep the old habits: copied extracts in every workspace, pipelines nobody owns, semantic models refreshed on a guess and capacity bills that surprise finance. The platform then looks like one more silo instead of a single analytics estate on OneLake. This Core Concept course trains engineers to build and run a Microsoft Fabric solution end to end, from workspaces and lakehouses through ingestion, Spark notebooks and warehouses to Direct Lake semantic models, security and capacity control. Participants leave with a Fabric Medallion Lakehouse Solution and Operating Runbook.

Course Objectives:

- Plan a Microsoft Fabric tenant layout of workspaces, capacities, domains and OneLake storage for a multi-team analytics estate
- Ingest source data into lakehouses and warehouses with Data Factory pipelines, Dataflows Gen2, shortcuts and mirroring
- Transform bronze data into silver and gold Delta tables using Fabric notebooks, Spark and T-SQL in the warehouse
- Publish Direct Lake semantic models and Power BI reports that read gold tables without scheduled import refresh
- Secure and govern Fabric items with workspace roles, row and column level security, sensitivity labels, endorsement, Git integration and deployment pipelines
- Operate the solution against capacity limits using the Capacity Metrics app, table maintenance and a documented runbook

Target Audience:

- Data engineers who build ingestion and transformation pipelines for shared analytics platforms
 - Analytics engineers who shape lakehouse and warehouse tables into business-ready models
 - BI developers moving Power BI datasets from import refresh onto OneLake and Direct Lake
 - Data platform administrators who manage workspaces, capacities, access and deployment
 - Database developers who migrate on-premises warehouses and marts into a SaaS analytics platform
-

Course Outline:

Day 1: Microsoft Fabric Platform Architecture and OneLake Foundations

- Fabric Workloads Map: Data Factory, Data Engineering, Data Warehouse, Real-Time Intelligence and Power BI
- OneLake Storage Model: One Copy, Delta Parquet Format and ADLS Gen2 Foundation
- Tenant, Capacity, Domain and Workspace Hierarchy Design
- OneLake Shortcuts to ADLS Gen2, Amazon S3 and Google Cloud Storage
- Current Analytics Estate Inventory and Fabric Adoption Gap Review

Day 2: Lakehouse, Warehouse and Medallion Architecture Design

- Medallion Layers: Bronze Raw, Silver Enriched and Gold Curated Zones
- Lakehouse Items: Files Area, Managed Tables and SQL Analytics Endpoint
- Fabric Data Warehouse: T-SQL Tables, Views, Stored Procedures and Cross-Database Queries
- Store Selection Matrix: Lakehouse, Warehouse, Eventhouse or Fabric SQL Database
- Layer-Per-Workspace Pattern and Star Schema Placement in the Gold Layer

Day 3: Ingestion and Transformation with Pipelines, Dataflows Gen2 and Notebooks

- Data Factory Pipelines: Copy Activity, Parameters, Schedules and Failure Paths
- Dataflows Gen2 Power Query Transformations with Lakehouse Destinations
- Mirroring Operational Databases into OneLake for Near Real-Time Replicas
- Fabric Notebooks with PySpark and Spark SQL for Bronze-to-Silver Cleansing
- Delta MERGE Upserts, Deduplication and Null Handling for Silver-to-Gold Loads

Day 4: Semantic Models, Real-Time Intelligence, Security and Capacity Control

- Direct Lake Semantic Models: Framing, Fallback to DirectQuery and Refresh Behaviour
- Real-Time Intelligence Overview: Eventstreams, Eventhouse, KQL Querysets and Real-Time Hub
- Workspace Roles, Item Permissions, OneLake Security and Row and Column Level Security
- Sensitivity Labels, Endorsement, OneLake Catalog and Lineage Impact Analysis
- Capacity Metrics App, Throttling, Smoothing, V-Order, OPTIMIZE and VACUUM Maintenance

Day 5: Lab Build and the Fabric Medallion Lakehouse Solution

- Retail Sales Lab: Pipeline and Notebook Load from Bronze to Gold Delta Tables
 - Utilities Asset Lab: Warehouse Star Schema and Direct Lake Power BI Report
 - Git Integration and Deployment Pipelines from Development to Production Workspace
 - Fabric Medallion Lakehouse Solution and Operating Runbook Assembly
 - Solution Walkthrough, Capacity Cost Review and Peer Feedback
-

Skills You Will Gain:

- Fabric Workspace Architecture
- OneLake Data Integration
- Lakehouse Engineering
- Pipeline Orchestration
- Spark Notebook Development
- Direct Lake Semantic Modelling
- Fabric Access Control
- Capacity Cost Management

Why Attend This Course:

- Return with a Fabric Medallion Lakehouse Solution and Operating Runbook built in the labs and ready to adapt to your own estate
- Replace duplicated extracts and ad hoc refreshes with one governed copy of data on OneLake
- Explain capacity consumption to finance and platform owners with evidence from the Capacity Metrics app
- Compare Fabric design decisions with engineers from retail, utilities, finance and public service organisations

Conclusion:

A Fabric estate earns trust when storage, pipelines, models and permissions are designed as one system rather than assembled item by item. The course moves from platform architecture and OneLake, through lakehouse, warehouse and medallion design, to ingestion with pipelines, Dataflows Gen2, mirroring and notebooks, and then to Direct Lake models, real-time options, security, governance and capacity control. The final day brings these together in two labs and produces a Fabric Medallion Lakehouse Solution and Operating Runbook for the workplace.
