

DP-700: Fabric Data Engineer

Engineer reliable batch and real-time data pipelines in Microsoft Fabric and prepare for the DP-700 Fabric Data Engineer exam.

DURATION

4 days · 32 hours

LEVEL

Intermediate

CURRICULUM

13 modules · 11 labs

DELIVERYLive online, Onsite,
Corporate

Course overview

Data engineering on Microsoft Fabric combines familiar tools — Spark, SQL and pipelines — with newer ones such as eventstreams, eventhouses and OneLake shortcuts. The data engineer's job is to choose the right tool for each workload, build pipelines that recover gracefully and keep costs and performance under control. This four-day course builds those skills on a realistic end-to-end scenario.

Day one covers implementing and managing an analytics solution: workspace settings, Spark configuration, security, lifecycle management with Git and deployment pipelines, and orchestration with pipelines and notebooks. Day two focuses on batch ingestion and transformation — loading data with pipelines and Dataflows Gen2, transforming it with PySpark and T-SQL, and handling incremental loads and slowly changing dimensions.

Day three is dedicated to real-time data: eventstreams, eventhouses, KQL querysets and Spark structured streaming, with windowing and late-arriving data. Day four covers monitoring and optimisation — the Monitoring hub, capacity metrics, error handling, Delta table maintenance and query performance — followed by exam preparation.

Every lab builds on the same scenario, so you finish with a working lakehouse, warehouse and real-time pipeline that you designed and tuned yourself.

The course is aligned to the DP-700 exam, which is the recommended Fabric path for data engineers following the retirement of DP-203. Fabric evolves quickly, so check the current outline on Microsoft Learn before booking.

Highlights

- 32 hours of live, hands-on sessions
- Batch and streaming pipelines built in one Fabric workspace
- PySpark, T-SQL and KQL used where each fits best
- Orchestration, monitoring and optimisation labs
- Recommended path for learners who planned DP-203

What's included

- Batch and streaming pipelines in one scenario
- Incremental loading and slowly changing dimensions
- Eventstreams, eventhouses and KQL labs
- Orchestration with pipelines and notebooks
- Delta table maintenance and performance tuning
- Monitoring hub and error-handling patterns
- 30 days of post-training doubt support

Who this course is for

Who should attend

- Data engineers moving to Microsoft Fabric
- ETL and data warehouse developers
- Learners who were preparing for the retired DP-203 exam
- Analytics engineers who build ingestion pipelines
- Developers working with streaming data

Prerequisites

- Working knowledge of SQL
- Basic Python or PySpark experience
- Understanding of data warehousing concepts (DP-900 level)

Detailed curriculum — 4 days, 32 hours, 13 modules, 11 hands-on labs

Day 1 Implementing and Managing an Analytics Solution

8 hours

1.1 Workspace and Spark configuration

- Workspace settings and capacities
- Domains and OneLake organisation
- Spark pools, environments and libraries
- Choosing between lakehouse, warehouse and eventhouse

LAB Configure a workspace with a custom Spark environment · 40 min

1.2 Security and lifecycle

- Workspace and item-level access
- Row-level, column-level and dynamic data masking in the warehouse
- OneLake data access roles
- Git integration and deployment pipelines

LAB Secure warehouse data with masking and row-level security · 45 min

1.3 Orchestration

- Pipelines vs notebooks for orchestration
- Schedules and event-based triggers
- Parameters and dynamic expressions
- Dependencies and retries

LAB Orchestrate notebooks from a parameterised pipeline · 45 min

Day 2 Batch Ingestion and Transformation

8 hours

2.1 Loading data

- Copy activity and connectors
- Shortcuts and mirroring
- Dataflows Gen2
- Full vs incremental loads

LAB Build an incremental load with a watermark · 50 min

2.2 Transforming with PySpark

- DataFrames and Spark SQL
- Writing Delta tables
- Joins, aggregations and window functions
- Handling schema drift

LAB Transform raw files into silver Delta tables · 50 min

2.3 Transforming with T-SQL

- Warehouse loading patterns
- Stored procedures in pipelines
- Slowly changing dimensions

LAB Implement a type 2 slowly changing dimension · 45 min

Day 3 Real-Time Data Engineering

8 hours

3.1 Eventstreams

- Sources and destinations
- Routing events to multiple stores
- In-stream transformations

LAB Stream device telemetry with an eventstream · 45 min

3.2 Eventhouses and KQL

- KQL databases and tables
- Materialised views and update policies
- KQL queries and functions
- Real-Time dashboards

LAB Query and visualise streaming data with KQL · 55 min

3.3 Spark structured streaming

- Streaming DataFrames
- Late-arriving data
- Windowing and watermarks
- Checkpoints and exactly-once processing

LAB Process a stream into Delta with structured streaming · 45 min

Day 4 Monitoring, Optimisation and Exam Readiness

8 hours

4.1 Monitoring

- Monitoring hub
- Alerts with Activator
- Capacity metrics
- Diagnosing failed runs

LAB Investigate and fix failing pipeline runs · 45 min

4.2 Error handling

- Retry and failure paths in pipelines
- Logging patterns in notebooks
- Data validation checks

4.3 Optimisation

- OPTIMIZE, VACUUM and V-Order
- Spark performance settings
- Partitioning and file sizes
- Warehouse and KQL query tuning


LAB Tune a slow Spark job and maintain Delta tables · 50 min

4.4 Exam readiness


- Exam structure and case studies
- Practice questions with review

Skills and tools

Skills you will build



Tools used in class



Certification exam

This course is aligned to the exam below and prepares you for it. The exam itself is booked and taken with the vendor.

Exam	Microsoft Certified: Fabric Data Engineer Associate
Exam code	DP-700
Duration	100 minutes
Passing score	700 / 1000 (scaled)
Format	Multiple choice, Multiple response, Drag and drop, Case study
Delivery	Online proctored, Test centre
Validity	1 year (free online renewal assessment)
Official page	https://learn.microsoft.com/credentials/certifications/exams/dp-700/

Exam skill areas covered

Monitor and optimize an analytics solution

Exam details are pending verification. Confirm the current exam code, format and fee on the official vendor page before booking.

Frequently asked questions

I studied for DP-203. Is DP-700 the right next step?

Yes. DP-203 has retired, and DP-700 is Microsoft's data engineering certification for Fabric. Many Spark and SQL skills carry over directly.

How much coding is involved?

Quite a lot: PySpark, T-SQL and KQL appear every day. Starter code is provided, so basic experience is enough.

Do I need my own Fabric capacity?

No. Lab access is arranged for the course, and the details are shared before Day 1.

What happens if Microsoft updates the DP-700 exam?

We map the curriculum to the current published skills outline and revise it when Microsoft changes the exam. Always check the latest outline on Microsoft Learn before booking.

Do you run weekend batches?

Yes. Weekend batches cover the same 32 hours across consecutive Saturdays and Sundays, with timings in IST that also suit learners in the Gulf.

Can you run a private batch for our team?

Yes. We run private online or onsite batches for corporate teams, and we can weight the labs towards your existing data sources and pipelines. Request a corporate quote with your team size and preferred dates.

Book a batch or talk to an advisor

Course page: aceskillsconsulting.com/course/dp-700-fabric-data-engineer · Email: info@aceskillsconsulting.com

Chennai office

Guindy, Chennai, Tamil Nadu

Delhi office

B-95, Old Post Office Street, Shakarpur
New Delhi 110092
Office: 98180 12523

Microsoft, Azure, Microsoft 365, Power BI and Microsoft Entra are trademarks of the Microsoft group of companies. All product names, logos, brands, trademarks and registered trademarks are the property of their respective owners. Ace Skills Consulting is an independent training provider; unless explicitly stated, it is not affiliated with, endorsed by or sponsored by any vendor mentioned on this site. Exam details are provided for guidance and may change — always confirm on the vendor's official website. Curriculum generated on 23 Oct 2026; topics and order may be adjusted by the trainer to suit each batch.