

Microsoft Power BI

Course Duration ~60 Hours · 8 Weeks · Live Online · Build-Along · Portfolio of Real Reports

Curriculum reviewed: July 2026 — the Power BI and Fabric stack moves fast; this syllabus tracks the current product

Data Modelling

DAX

Performance

Governance

Fabric & Copilot

PL-300

In short: Anyone can drag a field onto a canvas. The people who get hired build the data model underneath it — and can explain why the number is right. This is a live, build-along program that treats modelling and DAX as the real skill and visuals as the last mile: star schema design, filter context, performance (Import vs DirectQuery vs Direct Lake, VertiPaq, DAX Studio), enterprise delivery (row-level security, deployment pipelines, certified semantic models), and Microsoft Fabric with Copilot handled honestly — AI accelerates a good semantic model and amplifies the errors in a bad one, so modelling matters *more* now, not less. Across 9 modules you ship a star-schema sales model with a fiscal calendar, a Power Query messy-source rescue, an executive one-pager, a documented performance rescue, a secured dev/test/prod deployment, and a Direct Lake vs Import comparison — closing with a governed capstone defended in a stakeholder walkthrough. Built for fresh graduates from any degree background and for Excel-heavy finance, operations, marketing and supply chain professionals — no programming background required, and still credible to engineers. Aligned to PL-300.

Duration	8 weeks · ~60 hours · live + hands-on	Format	Live online · build-along · portfolio of real reports
Prerequisite	Comfort with spreadsheets. Basic SQL helpful but taught as needed. No programming background required.	Curriculum	9 modules, from the Power BI landscape to Fabric, Copilot and capstone
Certification	Aligned to PL-300: Microsoft Power BI Data Analyst; DP-600 (Fabric Analytics Engineer) is the natural next step	Ideal for	Fresh graduates from any background and Excel-heavy professionals moving into analytics

Course Highlights

- Every module leads with what you build and ship, not what you "learn about"
- Modelling first: star schema, filter context and a proper date table before any dashboard
- DAX to real depth — including debugging a measure that returns the wrong number
- Performance as a discipline: Import vs DirectQuery vs Direct Lake, VertiPaq, DAX Studio
- Enterprise delivery: RLS, deployment pipelines, certified semantic models, governance
- Fabric and Copilot handled honestly — AI amplifies your model, good or bad
- Beyond the dashboard: design for a decision, then defend your numbers in front of stakeholders

Hands-on Projects & Portfolio

- **Sales & Profitability Model** — star-schema-first, with YoY, margin and rolling-average measures and a proper fiscal calendar
- **Messy-Source Rescue** — inconsistent files and systems cleaned entirely in Power Query with folding preserved
- **Executive One-Pager** — designed for a specific decision, with drillthrough to detail
- **Performance Rescue** — cut a slow, bloated model's load time, with before/after evidence from Performance Analyzer and DAX Studio
- **Secured Deployment** — row-level security by region, an app with audiences, and a dev/test/prod deployment pipeline
- **Direct Lake vs Import** — a Direct Lake report over Fabric lakehouse data, compared against an Import model
- **Capstone** — end-to-end governed report plus a stakeholder presentation defending the numbers

Tools & Stack

Power BI Desktop · Power BI Service · Power Query / M · DAX · Microsoft Fabric · OneLake · Direct Lake · Copilot in Power BI · Tabular Editor · DAX Studio · Power BI Gateway · Excel · SQL Server · Azure SQL · SharePoint · Dataverse

Certification Alignment

- Syllabus mapped to **PL-300: Microsoft Power BI Data Analyst** across its published skill areas — preparing data, modelling data, visualising and analysing data, and managing and securing Power BI
- **DP-600 (Fabric Analytics Engineer)** is the natural next step toward enterprise-scale analytics engineering
- Presented as preparation and alignment — no official partnership or accreditation implied, and no exam outcome guaranteed

Career Outcomes

Power BI is the default analytics layer in organisations standardised on Microsoft 365, which makes the skill broadly employable across industries and functions rather than concentrated in tech companies — and it is one of the few technical roles reachable from a non-engineering background. The durable differentiator is modelling and DAX, not visual design: that is what separates a report builder from a BI developer. This course targets:

- Power BI Developer · Data Analyst · Business Intelligence Analyst · BI Developer
- Analytics Engineer · Reporting Analyst · Data Visualisation Specialist · Fabric Analytics Engineer (with DP-600)

Who This Course Is For

- Final-year students and fresh graduates targeting analyst roles, from any degree background
- Excel-heavy finance, operations, marketing, HR and supply chain professionals moving into analytics
- Software engineers, testers and support engineers pivoting toward data roles
- SQL developers and report writers modernising from legacy BI tools
- Team leads and managers who commission dashboards and need to judge whether they're any good
- Anyone preparing for PL-300

Module 1: The Power BI Landscape

5 Hours

YOU SHIP

- Your first published report — and an annotated map of the end-to-end path a number takes from source system to executive slide

TOPICS

- Desktop, Service and Mobile; how Power BI sits inside Microsoft Fabric
- Workspaces and capacities; semantic models as a shared asset
- The end-to-end path from source system to executive slide

Assessment: Quiz + Lab

Module 2: Data Preparation with Power Query

7 Hours

YOU SHIP

- A messy-source rescue: multiple inconsistent files and systems cleaned entirely in Power Query, with query folding preserved and refresh time measured

TOPICS

- Connecting to files, databases, cloud services and APIs; the M language where the UI runs out
- Shaping and cleaning; merges and appends; pivot/unpivot; parameters and functions
- Incremental refresh; query folding and why it decides your refresh time; dataflows for reusable prep

Assessment: Quiz + Lab

Module 3: Data Modelling — the Module That Matters Most

8 Hours

YOU SHIP

- A star schema rebuilt from a flat-table mess — facts, dimensions, a proper date table, deliberate cardinality — the model every later module builds on

TOPICS

- Star schema versus flat tables and why it isn't optional; fact and dimension design
- Relationships and cardinality; filter direction and ambiguity; role-playing dimensions
- A proper date table; many-to-many handling; composite models

Assessment: Quiz + Lab

Module 4: DAX Fundamentals

7 Hours

YOU SHIP

- A measure library for the sales model — readable, variable-structured, and correct under every filter combination you can throw at it

TOPICS

- Calculated columns versus measures and when each is wrong; row context versus filter context
- CALCULATE as the centre of the language; FILTER, ALL, VALUES, RELATED
- Iterators like SUMX and AVERAGEX; variables for readable measures

Assessment: Quiz + Lab

Module 5: Advanced DAX

7 Hours

YOU SHIP

- Time intelligence on a custom fiscal calendar — YoY, rolling averages, top-N with dynamic segmentation — plus a debugged fix for a measure that returned the wrong number

TOPICS

- Time intelligence: year-to-date, prior period, rolling averages, custom fiscal calendars
- Ranking and top-N; dynamic segmentation; what-if parameters; calculation groups
- Debugging a measure that returns the wrong number

Assessment: Quiz + Lab

Module 6: Visualisation & Report Design

7 Hours

YOU SHIP

- An executive one-page report designed for a specific decision — drillthrough to detail, accessible colour, mobile layout — then defended in a design critique

TOPICS

- Choosing visuals for the question being asked; layout and hierarchy
- Bookmarks, drillthrough, tooltips, field parameters, conditional formatting
- Accessibility and colour; mobile layouts; designing for a decision rather than for density

Assessment: Quiz + Lab

Module 7: Performance Optimisation

7 Hours

YOU SHIP

- A performance rescue: a slow, bloated model diagnosed and cut down, with before/after evidence from Performance Analyzer and DAX Studio

TOPICS

- Import vs DirectQuery vs Live vs Direct Lake; VertiPaq compression and cardinality
- Reducing model size; aggregations; Performance Analyzer; DAX Studio and Tabular Editor
- Capacity metrics; diagnosing why a report takes 40 seconds to load

Assessment: Quiz + Lab

Module 8: Deployment, Security & Governance

7 Hours

YOU SHIP

- A secured deployment: row-level security by region, an app with audiences, and a dev/test/prod deployment pipeline with scheduled refresh through a gateway

TOPICS

- Workspace roles and strategy; apps and audiences; scheduled and incremental refresh; gateways
- Row-level and object-level security; sensitivity labels; certified and promoted semantic models
- Deployment pipelines across dev/test/prod; version control; usage monitoring

Assessment: Quiz + Lab**Module 9: Fabric, Copilot & Capstone**

5 Hours

YOU SHIP

- A Direct Lake report over Fabric lakehouse data compared against an Import model — and your capstone: one end-to-end governed report from raw source to a stakeholder walkthrough where you defend the numbers

TOPICS

- OneLake and lakehouses; Direct Lake semantic models; connecting Power BI to Fabric workloads
- Copilot for report and measure assistance — and where it misleads you; semantic model descriptions that make AI features work
- Capstone build and stakeholder defense

Assessment: Capstone Project + Stakeholder Walkthrough