

Big Data / Analytics

Analytical SQL, dimensional modelling and dbt. You build tested, documented models that other teams can query without asking you first, and you keep the warehouse bill sane.

Foundational

4 phases

21 sessions

32 contact hours

Syllabus

4 phases, 21 sessions. Each phase ends in something you have built.

1

Questions and Answers

Answer a real business question in SQL without asking anyone for a cloud account.

5 sessions

\$72 alone

1.1 What Analytics Owns

The job in one sentence: a question in, a number someone acts on out, and why that is not the job of moving the data.

[FREE PREVIEW](#)

1.2 Getting Data In

Load CSV and Parquet, fix the column that arrived as text, and look at a dataset before trusting it.

1.3 SQL That Answers Something

SELECT, WHERE, GROUP BY, and reading your own result set with suspicion.

1.4 Joins and the Fan-Out

How one extra join doubles revenue, and how to catch it before anyone else does.

1.5 Windows and Cohorts

CTEs, window functions, retention, and building a hard query one readable step at a time.

By the end: **One SQL file that answers six questions about a raw dataset, run locally at zero cost.**

2

From Queries to Models

Turn a folder of queries into a project someone else can run.

5 sessions

\$72 alone

2.1 Dimensional Modeling

Kimball's four steps: pick the process, declare the grain, then facts and dimensions.

2.2 Your First dbt Project

Sources, refs, models, and a build order you never had to draw yourself.

2.3 Staging, Intermediate, Marts

The three layers, the naming rules, and why a 400-line query is not a model.

2.4 Materializations

View, table, and ephemeral: trading build time against query time on purpose.

2.5 Tests and Freshness

not_null, unique, relationships, and catching a broken source before anyone opens the dashboard.

By the end: **A layered dbt project on DuckDB that builds and tests itself from one command.**

3

Off the Laptop

Get the project onto a real warehouse and a real schedule without breaking it.

6 sessions

\$72 alone

3.1 Git for One Person

Branches, commits, pull requests, and getting yesterday's version back.

3.2 Into Snowflake

Accounts, virtual warehouses, storage and compute apart, micro-partitions, and a 30-day trial clock you start on purpose.

3.3 Same Project, Two Targets

Develop on DuckDB, deploy to Snowflake, and change one profile instead of every model.

3.4 Incremental Models and Snapshots

Rebuild only what changed, and keep SCD Type 2 history that does not quietly rewrite itself.

3.5 Docs and Lineage

Generated documentation, column-level lineage, and answering "where does this number come from" in ten seconds.

3.6 Deploy: CI and Schedules

Checks on every pull request, a nightly production run, and an alert when it fails at 3am.

By the end: **The same project running nightly on Snowflake, with tests gating every pull request.**

4

Trust, Cost, Consumers

Make the numbers agreed on, affordable, and actually used by someone.

5 sessions

\$72 alone

4.1 Metrics and the Semantic Layer

One definition of "active user", written once in YAML and reused by every tool that asks.

4.2 Questions Without SQL

Semantic views, Cortex Analyst, and MCP: a natural-language answer is only as good as the model underneath it.

4.3 Warehouse Cost

Credits, auto-suspend, adaptive warehouses, budgets, and a resource monitor that stops the bleeding.

4.4 Access and PII

Roles, grants, masking policies, row access, and a data request you are allowed to refuse.

4.5 Delivering Analytics

Dashboards as code, self-serve for people who will not write SQL, and the readout where someone disagrees with your number.

By the end: **A published metric with one named owner, one dashboard, and a monthly cost you can state.**

Tools you will use

Snowflake

dbt Fusion

DuckDB

Parquet

MetricFlow

Apache Ossie

Evidence

GitHub Actions

What you will build

1 Six Questions, One SQL File

A raw dataset answered on your laptop, zero cloud spend

2 Two Warehouses, One Project

Develop on DuckDB, ship to Snowflake, change one file

3 The Metric Nobody Argues About

One definition, one owner, one number

Registration

Phase 1: Questions and Answers	\$72
Phase 2: From Queries to Models	\$72
Phase 3: Off the Laptop	\$72
Phase 4: Trust, Cost, Consumers	\$72
Whole track, all 4 phases	\$240

Buying every phase separately costs \$288, so the whole track saves \$48.

No payment is taken online. Registering creates an order and payment instructions are sent with it. Access opens once payment is confirmed.