



Data Warehousing and Dimensional Modelling: ETL, ELT, Data Lakes and Lakehouse



Data Warehousing and Dimensional Modelling: ETL, ELT, Data Lakes and Lakehouse

Introduction:

Many organisations still report from spreadsheets stitched together from operational systems, so figures disagree, history is overwritten and every new question needs a new extract. Data warehousing solves this by integrating source data into a governed analytical store modelled for fast, consistent reporting. This Core Concept course equips practitioners to design star schemas, build reliable ETL and ELT pipelines, decide where a data lake or lakehouse fits, and secure and tune the platform for cost. Participants leave with a Data Warehouse Architecture and Dimensional Model Blueprint for a priority subject area.

Course Objectives:

- Assess analytical requirements and source systems to select a warehouse, data lake or lakehouse architecture that fits the organisation
- Design star and snowflake schemas with the correct grain, fact table type, conformed dimensions and surrogate keys
- Build ETL and ELT load routines that handle staging, incremental loads, change data capture and slowly changing dimensions
- Embed data quality tests, reconciliation checks and lineage capture inside warehouse pipelines
- Apply access control, row-level security, masking and encryption while tuning partitioning, clustering and compute for performance and cost
- Produce a Data Warehouse Architecture and Dimensional Model Blueprint ready for architecture board review

Target Audience:

- Managers and leads responsible for business intelligence platforms and enterprise reporting delivery
 - Data engineering leads accountable for ingestion, transformation and load pipelines
 - Solution and data architects responsible for analytical platform design decisions
 - Analytics managers whose teams depend on trusted historical and cross-functional data
 - Platform and cloud operations managers accountable for warehouse availability, security and running cost
-

Course Outline:

Day 1: Data Warehouse Foundations and the Analytical Current State

- Operational Versus Analytical Workloads: OLTP and OLAP Characteristics
- Inmon Corporate Information Factory Versus Kimball Dimensional Bus Architecture
- Data Warehouse, Data Lake and Data Lakehouse Comparison Matrix
- Source System Inventory and Reporting Pain Point Assessment
- Business Process Prioritisation With the Enterprise Bus Matrix

Day 2: Dimensional Modelling and Reference Architectures

- Kimball Four-Step Dimensional Design: Business Process, Grain, Dimensions and Facts
- Star Schema Versus Snowflake Schema Design Trade-Offs
- Transaction, Periodic Snapshot and Accumulating Snapshot Fact Tables
- Conformed Dimensions, Surrogate Keys and Role-Playing Date Dimensions
- Data Vault 2.0 Hubs, Links and Satellites as an Integration Layer

Day 3: Building ETL and ELT Pipelines

- ETL Versus ELT Pattern Selection and Staging Area Design
- Slowly Changing Dimension Type 1, Type 2 and Type 3 Load Logic
- Incremental Loading With Watermarks and Change Data Capture
- Medallion Architecture: Bronze, Silver and Gold Layer Responsibilities
- Pipeline Orchestration, Dependency Scheduling and Failure Recovery

Day 4: Lakehouse Platforms, Warehouse Security, Performance and Cost

- Open Table Formats Compared: Delta Lake, Apache Iceberg and Apache Hudi
- In-Pipeline Data Quality Tests, Reconciliation Counts and Column-Level Lineage
- Warehouse Security: Role-Based Access, Row-Level Security, Dynamic Masking and Encryption
- Partitioning, Clustering and Materialised Views for Query Performance
- Storage and Compute Separation, Workload Isolation and Cloud Cost Controls

Day 5: Lab Case Work and the Warehouse Blueprint

- Retail Sales Lab: Star Schema Build From Point-of-Sale Feeds
 - Healthcare Operations Lab: Late-Arriving Facts and Historical Dimension Tracking
 - Semantic Layer Lab: Shared Metrics and BI Tool Connection Over the Gold Layer
 - Data Warehouse Architecture and Dimensional Model Blueprint Drafting
 - Architecture Review Board Panel and Blueprint Defence
-

Skills You Will Gain:

- Dimensional Data Modelling
- Warehouse Architecture Selection
- ETL and ELT Pipeline Engineering
- Change Data Capture
- Lakehouse Table Format Evaluation
- Warehouse Access Security
- Query Performance Tuning
- Cloud Data Platform Cost Management

Why Attend This Course:

- Return with a Data Warehouse Architecture and Dimensional Model Blueprint for a subject area your organisation needs to report on
- Stop overwriting history and rebuilding extracts by loading dimensions and facts in a repeatable, auditable way
- Explain to sponsors when a warehouse, a data lake or a lakehouse is the right investment and what it will cost to run
- Test design choices in hands-on labs alongside data professionals from retail, healthcare, finance and public service

Conclusion:

A data warehouse earns trust when its model matches how the business measures itself and its pipelines load history the same way every night. This course moves from warehouse concepts and architecture choices, through dimensional modelling and ETL and ELT pipeline construction, to lakehouse table formats, in-pipeline quality checks, security, performance and cost. The final day applies these methods in retail and healthcare labs and produces a Data Warehouse Architecture and Dimensional Model Blueprint that participants take back to their platform team.
