

- COURSE CODE:** SQSDWD
- COURSE TITLE:** Data Warehouse Development for SQL Server 2005/2008/2008 R2
- AUDIENCE:** Data warehouse analysts, ETL specialists, data warehouse administrators.
- PREREQUISITES:** Completion of our course 'Data Warehousing Terms, Concepts & Architecture (DWHTCA)' and 6 months experience of Microsoft SQL Server or equivalent knowledge.
- DURATION:** 3 days
- LEVEL:** Intermediate
- FEATURES:** Data warehouse/data mart and multidimensional cube development using realistic scenarios.
- SUMMARY:** Designed to emphasise SQL Server features specific to business intelligence (BI) and data warehousing (DW), this scenario and hands-on presentation details the following:
- Dimensional modelling
 - Data mart design and load
 - Business Intelligence Development Studio (BIDS)
 - SQL Server Management Studio
 - Building dimension tables
 - Building fact tables
 - Developing dimensions
 - Hierarchies
 - ETL via SQL Server Integration Services (SSIS)
 - Security
 - SQL Server Analysis Services (SSAS)
 - SQL Server Reporting Services (SSRS)
 - Performance and good practices
- OBJECTIVES:** Upon completion of this course, the delegate should understand how to deal with BI/DW design issues as they relate to SQL Server 2005/2008/2008 R2.
- FORMAT:** Hands-on (about 60%).

1. REVIEW OF BI TERMS & CONCEPTS

- The data warehousing (DW) environment
- The DW computing context
- What is a data warehouse?
- What is a data mart?
- What is Business Intelligence (BI)?
- How do OLTP & OLAP differ?
- What is data mining?
- Operational vs. historical data
- Dimensional modelling
- Some dimensional modelling characteristics
- What is a star schema?
- What is a snowflake schema?
- What is metadata?

2. DATA MART DESIGN – INITIAL STEPS

- Determine business processes
- Determine grain
- Determine dimensions
- Determine facts

3. DATA MART DESIGN

- Fictitious organisation description
- Business process bus architecture
- Business questions to be answered and solutions to be met
- Star vs. snowflake schema
- Dimension design
- Hierarchy design
- Fact (measure) design

4. DATA MART CREATION: SQL SERVER MANAGEMENT STUDIO

- Database
- Dimension tables
- Fact tables
- Integrity and other relationships

5. ALTERNATIVE DATA MART CREATION: BIDS

- Analysis Services cube via Cube Wizard
- Data mart via Cube Definition

6. EXTRACT, TRANSFORM, LOAD (ETL) CONCEPTS

- ETL Overview
- Kimball's 38 ETL Subsystem
- ETL Staging Considerations
- Source-To-Destination Mapping Considerations
- Source System Monitoring Document
- Normalised (OLTP) vs. Dimensional Model
- Planning Source-To-ETL Processing
- How to Handle 'Inferred Members'
- Fact Table Properties & Overview

7. POPULATE DATA MART USING SQL SERVER INTEGRATION SERVICES (SSIS)


- Create project
- Use Control Flows
- Use Data Flows
- Use Source and Destination adaptors
- Use Connection Manager
- Use Source to Destination mapping
- Use Precedence Constraints
- Use Aggregate Transformation
- Load data into data mart

8. CREATE A SECOND FACT TABLE IN EXISTING DATA MART

- Rationale
- Use SQL Server Management Studio
- Use Properties window
- Establish referential integrity

9. POPULATE FACT TABLES: SSIS

- Use BIDS
- Create new project
- Use various Control Flows
- Use various Data Flows
- Establish Source and Destination
- Adjust and optimise SSIS data types
- Establish column mapping
- Use Sort Transformation
- Use Derived Column Transformation
- Use expressions



10. SQL SERVER ANALYSIS SERVICES (SSAS) & CUBES

- Rationale
- Extend existing data mart with OLAP cube
- Define measures for cube
- Determine granularity
- Define calculated measures
- How to format measures
- How to add measures to cubes
- How to relate dimensions to measure groups
- How to deploy OLAP cubes via BIDS
- How to deploy OLAP cubes via Analysis Services
- Deployment wizard
- Workshop

11. INTRODUCTION TO THE BUSINESS INTELLIGENCE WIZARD


- Rationale
- Key Performance Indicators (KPIs)
- Actions
- Partitions
- Aggregation

12. INTRODUCTION TO SQL SERVER REPORTING SERVICES (SSRS)

- Rationale
- Definition
- Architecture
- Report creation
- Report manager
- Management
- Ad-hoc reporting

13. PHYSICAL DESIGN CONSIDERATIONS

- Definition
- Set-up (e.g., estimated volumes, availability requirements)
- Configuration
- Software
- Indexes
- ETL design



14. DW GOOD PRACTICES & PERFORMANCE

- Hardware considerations
- Database considerations
- Table design
- Index design
- Partition design
- ETL design
- Schema design
- Data integrity design
- Security design
- Backup and recovery design
- Metadata design

15. INTEGRATION SERVICES TUNING & BEST PRACTICES

- How to optimise Data Flows
- How to optimise Data Sources
- How to optimise parsing
- How to optimise logging
- Index considerations
- How to optimise Lookup transformations
- How to optimise Fuzzy Lookup
- How to optimise Fuzzy Grouping
- Constraint best practices
- How to encourage parallel execution
- Buffer usage
- Execution trees
- Blocking vs. non-blocking transformations
- How to improve INSERT performance