

- COURSE CODE:** OWHPDP
- COURSE TITLE:** Oracle Physical Data Warehouse Design
- CURRENCY:** **10g & 11g**
- LEVEL:** Intermediate to advanced
- AUDIENCE:** Database administrators, system administrators and developers
- PREREQUISITES:** Completion of our course 'Data Warehouse Terms, Concepts & Architecture' (DWHTCA) or equivalent knowledge.
- DURATION:** 4 days (hands-on machine workshop)
- SUMMARY:** After discussing the components of the data warehouse and data warehousing, this course elaborates on the essential physical design and phases of an Oracle data warehouse (DW). Particular emphasis is placed on the following:
- DW partitioning
 - Index options (e.g., bitmaps)
 - Integrity and constraints
 - Materialized views (summary tables)
 - DW objects (dimensions, hierarchies, etc.)
 - The ETL process
 - Performance and good practices
- OBJECTIVES:** Upon completion of this course, the participant should be able to translate logical DW design specifications into an Oracle physical design model; and in turn, create the necessary scripts (e.g., DDL) to install a large scale partitioned DW, using, amongst other things, bitmap indexes.

1. REVIEW OF DATA WAREHOUSING (DW) TERMS AND CONCEPTS

- The DW environment
- 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
- What is a star schema?
- What is a snowflake schema?
- Normalization vs. denormalization
- What are hierarchies?
- What is dimensional modelling?
- What is the Data Warehouse Bus Architecture (DWB)?
- What are surrogate keys?
- What is Extract, Transform, Load (ETL)?
- What are Slowly Changing Dimensions (SCD)?
- What is Metadata?
- What are Materialized Views (MVs)?
- How does logical design differ from physical design?

2. ORACLE DATABASE ARCHITECTURE – A SUMMARY

- Memory structure
- Logical storage structure
- Physical storage structure
- Processes
- SQL tools and extensions
- Admin tools
- Real Application Cluster

3. LOGICAL → PHYSICAL DESIGN PROCESS

- DW database objects (e.g., partitions, dimensions, materialized views)
- Hardware and I/O considerations

4. DW PARTITIONING DESIGN

- Parallelism options
- Partitioning options
- Compression options
- Partitioning and join considerations
- NULL considerations
- Date data type considerations
- Composite (multi-key) considerations
- Local vs. global indexes
- Backup and recovery considerations

5. DW INDEX DESIGN OPTIONS

- B*tree index – how they work
- Bitmapped index – how they work
- Bitmapped join indexes – how they work
- NULL value considerations
- Partitioning indexes
- Local vs. global indexes
- Index-organised tables

6. JOIN OPTIONS

- Hash
- Lossless
- Nested loop
- Partition-wise
- Sort merge
- Star transformation

7. DW CONSTRAINT DESIGN

- Constraint types (unique, data cleanliness, optimisation, etc.)
- Constraint options (ENABLE NO VALIDATE, DISABLE NOVALIDATE)
- Constraints and partitioning
- Views and constraints

8. MATERIALIALIZED VIEWS (i.e., AGGREGATE/SUMMARY TABLES)

- When to consider
- Oracle's summary management feature
- How to create materialized views
- Physical storage considerations
- Populating (loading) options
- Refresh options
- Query rewrite design
- Query rewrite and integrity (e.g., ENFORCED, TRUSTED)
- How to use the Oracle Summary Advisor
- EXPLAIN MVIEW utility usage

9. DIMENSIONS

- Definition and role
- Hierarchies – types
- When to create dimensions
- How to create dimension
- Dimension with attributes
- Normalized dimensions
- Validating dimensions
- Altering dimensions
- Dimensions and constraints
- Deleting dimensions

10. THE ETL PROCESS

- Overview
- The ETL Process
- Extract
- Extraction Types
- Data Extraction Techniques
- Data Cleansing
- Data Cleansing Techniques
- Data Transformation
- Delivery
- Data Load Options
- Surrogate Keys
- ETL Sub-task Summary
- ETL Vendor Considerations
- ETL in the Database

11. ORACLE'S ETL OFFERINGS


- Overview
- Extraction from operational systems
- Changed data capture
- Data transformation (e.g., data integration & cleansing)
- Generating keys
- Loading the DW – options

12. ORACLE DW TOOLS

- Warehouse Builder
- Discover
- Reports
- OLAP & data mining

13. ORACLE DW PERFORMANCE CONSIDERATIONS

- ETL and Table Functions
- How to optimise bitmap and bitmap join indexes
- How to optimise bitmap star join transformations
- Placing Oracle objects (e.g., indexes) into separate buffers, enhancement of multiple block sizes
- Dynamic SGA feature
- How to use dynamic sampling
- Materialized views and FAST REFRESH
- Materialized view 'size' control
- Indexes
- Partitioning
- dbms_stats, OPTIMIZER_MODE
- Star vs. snowflake schemas



14. PARALLELISM & PERFORMANCE

- Introduction
- Parallelisable operation
- How parallelism works
- Parallelising SQL (DDL)
- Parallelising SQL (DML)
- Degree of parallelism
- Parameter setting for parallelism