

- COURSE CODE:** DWHPTW
- COURSE TITLE:** Data Warehousing Performance & Tuning Workshop
- PLATFORM:** DB2 for Linux/UNIX/Windows
- AUDIENCE:** DB2 BI and data warehouse developers, database administrators and system administrators
- PREREQUISITES:** Completion of our course 'DB2 Data Warehouse Fundamentals' (DWHFUN) or equivalent knowledge.
- DURATION:** 4 days
- SUMMARY:** This presentation and hands-on workshop emphasises how to optimise data warehouse design in both non-partitioned and partitioned environments.
Topics central to this presentation include:
- Giga- and Tera-byte table considerations
 - Partitioning and parallelism
 - I/O design
 - High performance and BI queries
 - DB/DBM configuration parameters
 - How to exploit the DB2 optimizer
 - Materialized Views/Materialized Query Tables (MQTs)
 - Multidimensional Clustering (MDC)
- OBJECTIVES:** Upon completion of this presentation, the participant should be able to design, optimise and implement large-scale data warehouses to meet a variety of BI (Business Intelligence).
- FORMAT:** Lecture and hands-on

1. REVIEW OF THE BASICS

- Data warehouse vs. data mart
- Data modelling options
- Data mining
- Metadata
- OLAP vs. OLTP
- ETL (Extract, Transform, Load)
- Operational vs. historical data

2. SYSTEM ENVIRONMENT CONSIDERATIONS

- I/O placement considerations
- Log considerations
- Connection considerations
- Buffer pool considerations
- Locking considerations
- Package cache considerations
- Catalog cache considerations
- Sort considerations
- Other memory considerations
- Miscellaneous considerations

3. PERFORMANCE & MONITORING TOOLS

- Primary performance parameters (e.g., DB & DBM CFG)
- Facilities (e.g., DB2 Design Advisor)
- Explain/ Visual Explain
- db2diag.log
- DB2 Admin. Notification log
- Health Monitor and Health Center

4. APPLICATION DESIGN

- Table design (e.g., sizing)
- Index design
- Table space design
- SQL design
- Locking and concurrency design
- Integrity constraints

5. DB2 UDB'S MATERIALIZED VIEWS

- Materialized view overview
- Materialized view CREATE considerations
- Materialized view maintenance considerations
- Loading base tables (LOAD utility)
- Materialized view ALTER considerations
- Materialized view DROP considerations
- Materialized view matching considerations
- Materialized view design considerations
- Materialized view tuning considerations
- Refresh optimization
- Materialized view limitations

6. DB2 UDB ENVIRONMENT CONSIDERATIONS

- Memory
- Buffer pools
- Package cache
- Sorts
- Logging
- I/O placement
- RAID considerations

7. COMMAND & UTILITY OPTIONS


- EXPORT
- IMPORT
- LOAD (e.g., MASS LOAD, multipartition)
- REORG
- RESTORE
- RUNSTATS

8. PARTITIONING & PERFORMANCE

- Partition scope (i.e., number)
- Partition groups
- Partitioning maps
- Partitioning keys

9. DB2 OPTIMIZER (COMPILER)

- Query performance
- Inter-partition communication
- Join strategies
- Parallelism



10. DB2 UDB'S STATISTICS, ANALYTIC, AND OLAP FUNCTIONS

- Statistics and analytic functions
- OLAP functions

11. MULTIDIMENSIONAL CLUSTERING (MDC)

- Advantages of MDC
- Logical and physical organization of data
- Indexes and performance
- Clusters and indexes
- Issues with single dimension indexes
- Multidimensional clustering (MDC) and its benefits
- Creating MDC tables
- How MDC works logically and physically
- Reading data from an MDC table
- Writing data to an MDC table
- MDC tuning