


DB2 UDB for LUW

Advanced SQL via DB2 Explain Facilities



- COURSE CODE:** DWUASE
- COURSE TITLE:** Advanced SQL via DB2 Explain Facilities
- AUDIENCE:** Application Developers, Application Programmers, Production DBAs, Database Administrators and Capacity Planners.
- PREREQUISITES:** Completion of our course 'DB2 SQL Essentials' (DWUSQE) or at least one year of SQL experience in the Windows, UNIX, Linux or OS/2 environment.
- DURATION:** 3 days
- SUMMARY:** This course focuses on tuning advanced SQL via DB2 Explain facilities in terms of the following DB2 UDB features:
- Isolation levels
 - Locking and concurrency
 - Join optimisation
 - Sequences and identity columns
 - Common Table Expressions
 - Recursive queries
 - Grouping sets
 - Materialized Query Tables (MQTs)
 - Index design and usage
 - SARGable vs. non-SARGable predicates
 - Elapsed, CPU and I/O times.
- OBJECTIVES:** Upon completion of this presentation, the participant should be able to exploit the DB2 Optimizer in deciding which of several SQL alternatives will provide reduced elapsed times, CPU and I/O.
- LAB:** About 65% hands-on, using DB2 Explain facilities.

1. DB2 UDB ENVIRONMENT & OBJECTS REVIEW

2. BASIC SQL REVIEW & SELF-ASSESSMENT

3. ADVANCED SQL FACILITIES

- Sequences and identity columns
- Common Table Expressions
- Recursive queries
- Grouping sets
- Materialized Query Tables (MQTs)

4. PREDICATE TYPES

- SARGable
- Non-SARGable
- Indexable
- Range delimiting
- Residual

5. I/O TYPES

- Table space scans
- Index scans
- Index screening
- Index only
- Multiple index access
- Clustered indexes
- Sequential prefetch
- List prefetch

6. SQL COMPILER & OPTIMIZER

- Compilation process
- RUNSTATS considerations
- Query rewrite
- Operation merging (e.g., Subquery Join)
- Operation movement(e.g., DISTINCT)
- Predicate translation (e.g., OR IN)
- Optimisation classes (0 through 9)

7. INTRODUCTION TO DB2 EXPLAIN

- Its purpose
- How to invoke
- Authorisation
- General syntax rules

8. INTRODUCTION TO DB2 EXPLAIN TABLES

- OBJECT
- PREDICATE
- STATEMENT
- STREAM
- ADVISE_INDEX
- ADVISE_WORKLOAD
- How to create
- Examples of use

9. EXPLAIN OUTPUT INTERPRETATION


- **dynexpln** utility
- RUNSTATS vs. no RUNSTATS
- Scenarios with and without indexes

10. EXPLAIN ANALYSIS & TUNING WORKSHOP

- Simple selects
- Clustered vs. non-clustered indexes
- BETWEEN, varying values
- Multiple index access and list prefetch
- Index negation (e.g., OPTIMIZE FOR n ROWS)
- UNIONS
- Joins
- Subqueries, including transformations
- IN vs. EXISTS
- NOT IN vs. EXISTS
- UNION vs. CASE
- DISTINCT vs. GROUP BY

11. JOIN CONCEPTS & STRATEGIES

- Nested loop join
- Merge scan join
- Hash join
- OUTER vs. INNER join
- Collocated join
- Broadcast outer-table joins
- Directed outer-table joins
- Directed inner-table & outer-table joins
- Broadcast inner-table joins
- Directed inner-table joins
- Table queues
- Parallel joins
- EXPLAIN workshop



12. INTRODUCTION TO VISUAL EXPLAIN

- Concepts and definition
- How to invoke
- How to interpret output
- Importance of RUNSTATS
- Workshop

13. LOCKING & CONCURRENCY

- Purpose
- Table space locks
- Table locks
- Row locks
- Isolation levels
- Specifying and ALTERing
- Lock conversion/escalation
- Deadlock and timeout
- LOCK TABLE statement
- CURSOR...WITH HOLD
- CLOSE CURSOR WITH RELEASE
- COMMIT/ROLLBACK

14. EFFICIENT APPLICATION & SQL GUIDELINES

- Number of SQL calls
- Number of rows searched
- Number of columns retrieved
- SARGable vs. non-SARGable
- Static vs. dynamic SQL
- Joins
- Local vs. Join predicate
- Nested table expressions
- CASE vs. UNION
- Singleton select vs. Cursor
- OPTIMIZE FOR n ROWS
- OPTIMIZE FOR FIRST n ROWS
- Controlling sorts