



COURSE CODE: DZOASE

COURSE TITLE: Advanced SQL via DB2 EXPLAIN

- Using **EXPLAIN tables** (i.e., PLAN_TABLE/
DSN_STATEMNT_TABLE)

AUDIENCE: Application Developers, Application Programmers, Production DBAs, Database Administrators and Capacity Planners.

PREREQUISITES: Completion of course 'DB2 SQL Essentials' (DZOSQE) or equivalent knowledge.

DURATION: 4 days

SUMMARY: After reviewing the usage and the scope and limits of SQL operations such as:

- INNER JOINS
- OUTER JOINS
- UNIONS
- CASE expressions
- Subqueries
- Subselects
- Table expressions
- Row expressions
- Column/scalar functions
- Date/time arithmetic
- Cursor operations
- Locking & concurrency
- Index design & usage

this course focuses on using **DB2 EXPLAIN** (PLAN_TABLE and DSN_STATEMNT_TABLE) to optimise SQL statements in terms of 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, incorporating DB2 EXPLAIN

1. DB2 OBJECTS – REVIEW

- Storage groups
- Databases
- Table spaces
- Tables
- Indexes
- Views
- Aliases vs. synonyms
- MetaData via SYSIBM tables

2. DATA INTEGRITY

- Referential integrity
- Check constraints
- Unique constraints
- Primary vs. foreign keys

3. THE DB2 ENVIRONMENT

- Address spaces
- Database/SQL services
- Catalog services
- Logging/recovery services
- Locking/concurrency services
- Caching/buffering services

4. ZPARMs -- APPLICATION SPECIFIC

- EDM pool
- SORT pool
- RID pool
- DSMAX

5. THE DB2 OPTIMIZER

- Definition (e.g., its inputs)
- Function during BIND/REBIND
- Use of filter factors
- Role of RUNSTATS
- I/O time estimation
- Processor time estimation
- Access path selection
- Fooling the Optimizer

6. PLAN_TABLE/DSN_STATEMNT_TABLE – OUTPUTS (HOW TO INTERPRET)

- Function
- How to create
- How to populate (e.g., BIND)
- Index access
- Joins
- Sorts
- Table space scans
- Prefetch
- UNIONS
- Subqueries
- Locking
- View materialisation
- Parallelism
- CPU time estimates (i.e., Milliseconds)

7. 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

8. OPTIMISATION HINTS EXAMPLES

- When to use
- How to enable
- How to use
- How to locate in PLAN_TABLE

9. INDEX DESIGN & USAGE

- Index structure
- Advantages/disadvantages
- Clustered index criteria
- Composite indexes
- Index lookaside
- When to reorganise

10. I/O TYPES

- Table space scans
- Prefetch
- Synchronous vs. asynchronous
- ROWID

11. PREDICATE TYPES

- Definitions
- Indexable vs. non-indexable
- Stage 1 vs. Stage 2
- Stage 3 (i.e., host program implemented)
- Predicate evaluation order

12. DATE-TIME ARITHMETIC

- YEAR/MONTH calculations
- Last day of month calculations
- TIME calculations
- Subtracting DATE/TIME values
- DAYS function
- Converting DATE/TIME types

13. TEMPORARY TABLE USAGE GUIDELINES

- When to consider
- CREATE vs. DECLARED temporary tables
- DBA responsibility
- Application responsibility
- Examples of use

14. PLANS vs. PACKAGES

- One large plan
- Many packages
- Collections
- Version controls
- ACQUIRE & RELEASE parameters
- CURRENTDATA options
- ISOLATION LEVEL options
- CICS vs. Batch

15. BATCH APPLICATION PERFORMANCE

- Index lookaside
- Prefetch
- Utilities (e.g., application INSERTs vs. LOAD)
- Caching data in working storage
- How to reduce SQL calls
- Checkpoint/restart

16. ONLINE (e.g., CICS) PERFORMANCE

- Concurrency
- Thread reuse
- Controlling 'hot spots'
- Indexes
- DB2CONN, DB2ENTRY & DB2TRAN (or RCT) and performance

17. LOCKING & CONCURRENCY

- Reasons for locking
- Commits
- Page/row locks
- Table space locks
- Partition locks
- Isolation levels
- Lock avoidance
- BIND/REBIND parameters

18. EFFICIENT SQL GUIDELINES SUMMARY

- Number of SQL calls
- Number of rows searched
- Number of columns retrieved
- Stage 1 vs. Stage 2 predicates
- 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

19. WHAT'S NEW IN DB2 V9 (FOR DEVELOPERS)