

Db2 12 for z/OS SQL Performance and Tuning

OEM: IBM • Code: CV964G

COURSE DESCRIPTION

This course is designed to teach the students how to prevent SQL performance problems and how to improve the performance of existing SQL.

AUDIENCE PROFILE

- This course is for Db2 12 for z/OS application developers, Db2 12 for z/OS DBAs, and anyone else with a responsibility for SQL performance and tuning in a Db2 12 for z/OS environment.

PREREQUISITES

- Familiarity with SQL
- Familiarity with Db2 12 for z/OS
- Familiarity with Db2 12 for z/OS application programming

COURSE OBJECTIVES

By the end of this course, participants will be able to:

- Understand and design better indexes
- Determine how to work with the optimizer (avoid pitfalls, provide guidance)
- Optimize multi-table access
- Work with subqueries
- Avoid locking problems
- Use accounting traces and other tools to locate performance problems in existing SQL

COURSE OUTLINE

Module 1: Introduction to SQL performance and tuning

- Performance issues
- Simple example
- Visualizing the problem

- Summary

Module 2: Performance analysis tools

- Components of response time
- Time estimates with VQUBE3
- SQL EXPLAIN
- The accounting trace
- The bubble chart
- Performance thresholds

Module 3: Index basics

- Indexes
- Index structure
- Estimating index I/Os
- Clustering index
- Index page splits

Module 4: Access paths

- Classification
- Matching versus Screening
- Variations
- Hash access
- Prefetch
- Caveat

Module 5: More on indexes

- Include index
- Index on expression
- Random index
- Partitioned and partitioning, NPSI and DPSI
- Page range screening
- Features and limitations

Module 6: Tuning methodology and index cost

- Methodology
- Index cost: Disk space
- Index cost: Maintenance
- Utilities and indexes
- Modifying and creating indexes
- Avoiding sorts

Module 7: Index design

- Approach
- Designing indexes

Module 8: Advanced access paths

- Prefetch
- List prefetch
- Multiple index access
- Runtime adaptive index

Module 9: Multiple table access

- Join methods
- Join types
- Designing indexes for joins
- Predicting table order

Module 10: Subqueries

- Correlated subqueries
- Non-correlated subqueries
- ORDER BY and FETCH FIRST with subqueries
- Global query optimization
- Virtual tables
- Explain for subqueries

Module 11: Set operations (optional)

- UNION, EXCEPT, and INTERSECT Rules
- More about the set operators
- UNION ALL performance improvements

Module 12: Table design (optional)

- Number of tables
- Clustering sequence
- Denormalization
- Materialized query tables (MQTs)
- Temporal tables
- Archive enabled tables

Module 13: Working with the optimizer

- Indexable versus non-indexable predicates

- Boolean versus non-Boolean predicates
- Stage 1 versus stage 2
- Filter factors
- Helping the optimizer
- Pagination

Module 14: Locking issues

- The ACID test
- Reasons for serialization
- Serialization mechanisms
- Transaction locking
- Lock promotion, escalation, and avoidance

Module 15: More locking issues (optional)

- Skip locked data
- Currently committed data
- Optimistic locking
- Hot spots
- Application design
- Analyzing lock waits

Module 16: Massive batch (optional)

- Batch performance issues
- Buffer pool operations
- Improving performance
- Benefit analysis
- Massive deletes