

Db2 12 for z/OS SQL Performance and Tuning

Duration: 3 Days **Course Code: CV964G** **Delivery Method: Virtual Learning**

Overview:

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

Virtual Learning

This interactive training can be taken from any location, your office or home and is delivered by a trainer. This training does not have any delegates in the class with the instructor, since all delegates are virtually connected. Virtual delegates do not travel to this course, Global Knowledge will send you all the information needed before the start of the course and you can test the logins.

Target Audience:

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.

Objectives:

- After completing this course, students 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 and more

Prerequisites:

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

Content:

■ Introduction to SQL performance and tuning	■ Features and limitations Tuning methodology and index cost	■ Number of tables
■ Performance issues	■ Methodology	■ Clustering sequence Denormalization
■ Simple example	■ Index cost: Disk space	■ Materialized query tables (MQTs)
■ Visualizing the problem	■ Index cost: Maintenance	■ Temporal tables
■ Summary Performance analysis tools	■ Utilities and indexes	■ Archive enabled tables Working with the optimizer
■ Components of response time	■ Modifying and creating indexes	■ Indexable versus non-indexable predicates
■ Time estimates with VQUBE3	■ Avoiding sorts Index design	■ Boolean versus non-Boolean predicates
■ SQL EXPLAIN	■ Approach	■ Stage 1 versus stage 2
■ The accounting trace	■ Designing indexes Advanced access paths	■ Filter factors
■ The bubble chart	■ Prefetch	■ Helping the optimizer
■ Performance thresholds Index basics	■ List prefetch	■ Pagination Locking issues
■ Indexes	■ Multiple index access	■ The ACID test
■ Index structure	■ Runtime adaptive index Multiple table access	■ Reasons for serialization
■ Estimating index I/Os	■ Join methods	■ Serialization mechanisms
■ Clustering index	■ Join types	■ Transaction locking
■ Index page splits Access paths	■ Designing indexes for joins	■ Lock promotion, escalation, and avoidance
■ Classification	■ Predicting table order Subqueries	■ More locking issues (optional)
■ Matching versus Screening	■ Correlated subqueries	■ Skip locked data
■ Variations	■ Non-correlated subqueries	■ Currently committed data
■ Hash access	■ ORDER BY and FETCH FIRST with subqueries	■ Optimistic locking
■ Prefetch	■ Global query optimization	■ Hot spots
■ Caveat More on indexes	■ Virtual tables	■ Application design
■ Include index	■ Explain for subqueries Set operations (optional)	■ Analyzing lock waits Massive batch (optional)
■ Index on expression	■ UNION, EXCEPT, and INTERSECT	■ Batch performance issues
■ Random index	■ Rules	■ Buffer pool operations
■ Partitioned and partitioning, NPSI and DPSI	■ More about the set operators	■ Improving performance
■ Page range screening	■ UNION ALL performance improvements	■ Benefit analysis
	■ Table design (optional)	■ Massive deletes

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/