

## Db2 12 for z/OS SQL Performance and Tuning

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

### 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:

<ul style="list-style-type: none"><li>■ Introduction to SQL performance and tuning</li><li>■ Performance issues</li><li>■ Simple example</li><li>■ Visualizing the problem</li><li>■ Summary Performance analysis tools</li><li>■ Components of response time</li><li>■ Time estimates with VQUBE3</li><li>■ SQL EXPLAIN</li><li>■ The accounting trace</li><li>■ The bubble chart</li><li>■ Performance thresholds Index basics</li><li>■ Indexes</li><li>■ Index structure</li><li>■ Estimating index I/Os</li><li>■ Clustering index</li><li>■ Index page splits Access paths</li><li>■ Classification</li><li>■ Matching versus Screening</li><li>■ Variations</li><li>■ Hash access</li><li>■ Prefetch</li><li>■ Caveat More on indexes</li><li>■ Include index</li><li>■ Index on expression</li><li>■ Random index</li><li>■ Partitioned and partitioning, NPSI and DPSI</li><li>■ Page range screening</li></ul>	<ul style="list-style-type: none"><li>■ Features and limitations Tuning methodology and index cost</li><li>■ Methodology</li><li>■ Index cost: Disk space</li><li>■ Index cost: Maintenance</li><li>■ Utilities and indexes</li><li>■ Modifying and creating indexes</li><li>■ Avoiding sorts Index design</li><li>■ Approach</li><li>■ Designing indexes Advanced access paths</li><li>■ Prefetch</li><li>■ List prefetch</li><li>■ Multiple index access</li><li>■ Runtime adaptive index Multiple table access</li><li>■ Join methods</li><li>■ Join types</li><li>■ Designing indexes for joins</li><li>■ Predicting table order Subqueries</li><li>■ Correlated subqueries</li><li>■ Non-correlated subqueries</li><li>■ ORDER BY and FETCH FIRST with subqueries</li><li>■ Global query optimization</li><li>■ Virtual tables</li><li>■ Explain for subqueries Set operations (optional)</li><li>■ UNION, EXCEPT, and INTERSECT</li><li>■ Rules</li><li>■ More about the set operators</li><li>■ UNION ALL performance improvements</li><li>■ Table design (optional)</li></ul>	<ul style="list-style-type: none"><li>■ Number of tables</li><li>■ Clustering sequence Denormalization</li><li>■ Materialized query tables (MQTs)</li><li>■ Temporal tables</li><li>■ Archive enabled tables Working with the optimizer</li><li>■ Indexable versus non-indexable predicates</li><li>■ Boolean versus non-Boolean predicates</li><li>■ Stage 1 versus stage 2</li><li>■ Filter factors</li><li>■ Helping the optimizer</li><li>■ Pagination Locking issues</li><li>■ The ACID test</li><li>■ Reasons for serialization</li><li>■ Serialization mechanisms</li><li>■ Transaction locking</li><li>■ Lock promotion, escalation, and avoidance</li><li>■ More locking issues (optional)</li><li>■ Skip locked data</li><li>■ Currently committed data</li><li>■ Optimistic locking</li><li>■ Hot spots</li><li>■ Application design</li><li>■ Analyzing lock waits Massive batch (optional)</li><li>■ Batch performance issues</li><li>■ Buffer pool operations</li><li>■ Improving performance</li><li>■ Benefit analysis</li><li>■ Massive deletes</li></ul>
--	---	--

## Further Information:

For More information, or to book your course, please call us on 00 966 92000 9278

[training@globalknowledge.com.sa](mailto:training@globalknowledge.com.sa)

[www.globalknowledge.com/en-sa/](http://www.globalknowledge.com/en-sa/)

Global Knowledge - KSA, 393 Al-Uroubah Road, Al Worood, Riyadh 3140, Saudi Arabia