Db2 12 for z/OS SQL Performance and Tuning

Duration: 3 Days **Course Code: CV964G**

Overview:

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

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

Further Information:

For More information, or to book your course, please call us on Head Office 01189 123456 / Northern Office 0113 242 5931

info@globalknowledge.co.uk

www.globalknowledge.com/en-gb/

Global Knowledge, Mulberry Business Park, Fishponds Road, Wokingham Berkshire RG41 2GY UK