Db2 12 for z/OS SQL Performance and Tuning (all optional themes included)

Trainings-ID: **TD-CV964Gspecial**

Zum Seminar →

Das nehmen Sie mit

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

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

Programm

- Introduction to SQL performance and tuning
- Performance issues
- Simple example
- Visualizing the problem
- SummaryPerformance analysis tools
- Components of response time
- Time estimates with VQUBE3
- SQL EXPLAIN

- The accounting trace
- The bubble chart
- Performance thresholdsIndex basics
- Indexes
- Index structure
- Estimating index I/Os
- Clustering index
- Index page splitsAccess paths
- Classification
- Matching versus Screening
- Variations
- Hash access
- Prefetch
- CaveatMore on indexes
- Include index
- Index on expression
- Random index
- Partitioned and partitioning, NPSI and DPSI
- Page range screening
- Features and limitationsTuning methodology and index cost
- Methodology
- Index cost: Disk space
- Index cost: Maintenance
- Utilities and indexes
- Modifying and creating indexes
- Avoiding sortsIndex design
- Approach
- Designing indexesAdvanced access paths
- Prefetch
- List prefetch
- Multiple index access
- Runtime adaptive indexMultiple table access
- Join methods
- Join types

- Designing indexes for joins• Predicting table orderSubqueries
- Correlated subqueries
- Non-correlated subqueries
- ORDER BY and FETCH FIRST with subqueries
- Global query optimization
- Virtual tables
- Explain for subqueriesSet operations (optional)
- UNION, EXCEPT, and INTERSECT
- Rules
- More about the set operators
- UNION ALL performance improvementsTable design (optional)
- Number of tables
- Clustering sequence
- Denormalization
- Materialized query tables (MQTs)
- Temporal tables
- Archive enabled tablesWorking with the optimize
- ndexable versus non-indexable predicates
- Boolean versus non-Boolean predicates
- Stage 1 versus stage 2
- Filter factor
- Helping the optimizer
- PaginationLocking issues
- The ACID test
- Reasons for serialization
- Serialization mechanisms
- Transaction locking
- Lock promotion, escalation, and avoidanceMore locking issues (optional)
- Skip locked data
- Currently committed data
- Optimistic locking
- Hot spots
- Application design
- Analyzing lock waitsMassive batch (optional)

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

Termine & Optionen