



TRAINING OFFERING

You can reach us at:

Arrow ECS, Nidderdale House, Beckwith Knowle, Harrogate, HG3 1SA

Email: educationteam.ecs.uk@arrow.com

Phone: 0870 251 1000



Db2 12 for z/OS SQL Performance and Tuning

| CODE: | LENGTH: | PRICE: |
|--------|-------------------|-----------|
| CV964G | 24 Hours (3 days) | £1,950.00 |

Description

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

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

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.

Prerequisites

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

Programme

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

Session Dates

| Date | Location | Time Zone | Language | Type | Guaranteed | PRICE |
|-------------|-------------------|-----------|----------|-----------------------|------------|-----------|
| 11 Nov 2024 | Virtual Classroom | GMT | English | Instructor Led Online | | £1,950.00 |

Additional Information

[This training is also available as onsite training. Please contact us to find out more.](#)