ProTech
protechtraining.com

"Charting the Course ...

... to Your Success!"

DB2 UDB for z/OS: Visual Explain for DB2 z/OS & SQL Tuning Course Summary

Description
This course focuses on using Visual Explain V8 for DB2 z/OS to optimize SQL statements in terms of elapsed, CPU and I/O times.

Objectives
At the end of this course, students will be able to:

- Exploit the DB2 Optimizer in deciding which of several SQL alternatives will provide reduced elapsed times, CPU and I/O.

Topics

- DB2 Objects
- Data Integrity
- The DB2 Environment
- ZPARMs
- Index Design & Usage
- DB2's Access Types
- Predicate Types
- Date Time Arithmetic
- Temporary Table Usage Guidelines
- Plans vs. Packages
- Batch Application Performance
- Locking & Concurrency
- Efficient SQL Guidelines Summary
- The DB2 Optimizer
- Visual Explain (VE) Introduction
- Visual Explain (VE) Terms & Concepts
- How to Tune SQL via VE
- What's New in DB2 V9 (for Developers)

Audience

This course is designed for Application Developers, Application Programmers, Production DBAs, Database Administrators and Capacity Planners.

Prerequisites

Students should have completed DB2 SQL Essentials' or equivalent knowledge.

Duration

Four days

ProTech
protechtraining.com

"Charting the Course ...

... to Your Success!"

DB2 UDB for z/OS: Visual Explain for DB2 z/OS & SQL Tuning

Course Outline

I. DB2 Objects Overview
- A. Storage groups
- B. Databases
- C. Table spaces
- D. Tables
- E. Indexes
- F. Views
- G. Aliases vs. synonyms
- H. MetaData via SYSIBM tables

II. Data Integrity Overview
- A. Referential integrity
- B. Check constraints
- C. Unique constraints
- D. Primary vs. foreign keys

III. The DB2 Environment Overview
- A. Address spaces
- B. Database/SQL services
- C. Catalog services
- D. Logging/recovery services
- E. Locking/concurrency services
- F. Caching/buffering services

IV. ZPARMs Application Specific
- A. EDM pool
- B. SORT pool
- C. RID pool
- D. DSMAX

V. Index Design & Usage
- A. Index structure
- B. Advantages/disadvantages
- C. Clustered index criteria
- D. Composite indexes
- E. Index look aside
- F. When to reorganize

VI. DB2's Access Types
- A. Table space scan
- B. Sequential prefetch
- C. List prefetch
- D. Dynamic prefetch
- E. Matching index scan
- F. Non-matching index scan
- G. Index screening
- H. Multiple index access
- I. Index look-aside
- J. Direct row access

VII. Predicate Types
- A. Definitions
- B. Index able vs. non-index able
- C. Stage 1 vs. Stage 2
- D. Stage 3' (i.e., host program implemented)
- E. Predicate evaluation order

VIII. Date Time Arithmetic
- A. YEAR/MONTH calculations
- B. Last day of month calculations
- C. TIME calculations
- D. Subtracting DATE/TIME values
- E. DAYS function
- F. Converting DATE/TIME types

IX. Temporary Table Usage Guidelines
- A. When to consider
- B. CREATE vs. DECLARED temporary tables
- C. DBA responsibility
- D. Application responsibility
- E. Examples of use
- F. Workshop

X. Plans vs. Packages
- A. One large plan
- B. Many packages
- C. Collections
- D. Version controls
- E. ACQUIRE & RELEASE parameters
- F. CURRENTDATA options
- G. ISOLATION LEVEL options
- H. CICS vs. Batch

ProTech
protechtraining.com

"Charting the Course ...

... to Your Success!"

DB2 UDB for z/OS: Visual Explain for DB2 z/OS & SQL Tuning

Course Outline (cont'd)

XI. Batch Application Performance
- A. Index look aside
- B. Prefetch
- C. Utilities (e.g., application INSERTs vs. LOAD)
- D. Caching data in working storage
- E. How to reduce SQL calls
- F. Checkpoint/restart

XII. Locking & Concurrency
- A. Reasons for locking
- B. Commits
- C. Page/row locks
- D. Table space locks
- E. Partition locks
- F. Isolation levels
- G. Lock avoidance
- H. BIND/REBIND parameters optimization

XIII. Efficient SQL Guidelines Summary
- A. Number of SQL calls
- B. Number of rows searched
- C. Number of columns retrieved
- D. Stage 1 vs. Stage 2 predicates
- E. Static vs. dynamic SQL
- F. Joins
- G. Local vs. join predicate
- H. Nested table expressions
- I. CASE vs. UNION
- J. Singleton select vs. cursor
- K. OPTIMIZE FOR n ROWS
- L. OPTIMIZE FOR FIRST n ROWS
- M. Controlling sorts
- N. Workshop

XIV. TheDB2 Optimizer
- A. Definition (e.g., its inputs)
- B. Function during BIND/REBIND
- C. Use of filter factors
- D. Role of RUNSTATS
- E. I/O time estimation
- F. Processor time estimation
- G. Access path selection
- H. Fooling the Optimizer

XV. Visual Explain (VE) – Introduction
- A. Purpose, scope and limits
- B. Installation
- C. Control Center plug-ins
- D. Prerequisites
- E. Inter-DB2 subsystem dependencies
- F. Statistics Advisor
- G. Tune SQL
- H. List static SQL
- I. View External Explain tables
- J. List cache statements
- K. Detail ZPARMs
- L. Stored procedures used by VE
- M. How to invoke
- N. Workshop

XVI. Visual Explain (VE) Terms & Concepts
- A. Access plan
- B. Query blocks
- C. Nodes
- D. Node plane
- E. Diagrams
- F. Canvas
- G. Workshop

XVII. How to Tune SQL via VE
- A. SQL statement source
- B. VE's menu and tools
- C. How to read/interpret VE's access plan graph
- D. How to tune predicate processing and filtering
- E. How to tune joins
- F. How to tune sorts
- G. How to analyse parallel queries
- H. How to generate reports
- I. How to use the Statistics Advisor
- J. Detailed workshop

XVIII. What's New in DB@ V9 (for Developers)