

DB2 Performance and Tuning for LINUX, UNIX and Windows Course Summary

Description

This course is designed to review many aspects of performance and tuning of DB2 version 9.7, for Linux, UNIX and Windows (LUW). DB2 commands and tools created by IBM for DB2 will be used to support workshop activities. Performance and tuning varies with type of user to be supported. We will review this topic for support of SQL users, programmers, data designers and database administrators.

Objectives

At the completion of this course, the student will be able to:

- Code SQL to improve performance
- Use DB2 EXPLAIN tools to evaluate performance expectations
- Design application programs to improve performance and review with EXPLAIN tools
- Use the tool DB2BATCH to benchmark applications
- Use various tools and snapshot functions and administrative views to evaluate DB2 performance and operational behaviour of applications
- Improve performance through parallelism
- Review data placement to improve performance
- Understand row compression and range partitioning on tables to improve performance
- Consider Multi=Dimensional clustering as a design option, by analyzing data and application needs
- Tune registry variables and configuration parameters to effect performance
- Use DB2 monitor tools to assist in performance tuning
- Use performance-improving options for DB2 utilities
- Configure DB2 for self tuning
- Use advanced features of utilities including support for identity columns
- Set for automatic maintenance.
- Use and analyse benefits of advanced table design like row compression.

Topics

- Creating tables and tablespaces
- Advanced SQL techniques
- SQL performance
- Data Access and EXPLAIN
- The DB2 Optimizer
- Multi-Dimensional Clustering
- Advanced utilities and table design options, including row compression and range partitioning
- Configuring for self tuning of DB2 and automatic maintenance
- Event and snapshot Monitors
- DB2 Operational Model
- The Health and Activity Monitors
- DB2 Configuration Parameters

Audience

This course is designed for Database Administrators and those senior support persons who are responsible for tuning DB2 for performance.

Prerequisites

Students should have Relevant Operating Systems experience (LINUX, UNIX or Windows) along with knowledge of relational database theory and SQL. Previous experience of using DB2 in a support role that includes some database administration is also necessary.

Duration

Three days

Due to the nature of this material, this document refers to numerous hardware and software products by their trade names. References to other companies and their products are for informational purposes only, and all trademarks are the properties of their respective companies. It is not the intent of ProTech Professional Technical Services, Inc. to use any of these names generically

DB2 System Performance and Tuning for LINUX, UNIX and Windows

Course Outline

- I. Creating tables and tablespaces**
 - A. Review concepts of SMS and DMS tablespaces
 - B. Show table options like identity columns and generated columns
 - C. Introduction to data row compression
 - D. Basic design using data range partitioning
- II. Advanced SQL techniques**
 - A. How to use different outer join techniques
 - B. Using OLAP functions
 - C. Review Nested and common table expressions
 - D. Using recursion with common table expressions
- III. SQL performance**
 - A. Hints and tips for effectively using SQL
- IV. Data Access and EXPLAIN**
 - A. Review of Data access techniques
 - B. How to use various explain tools and interpret results
- V. The DB2 Optimizer**
 - A. Review of strategy to improve optimization strategy
 - B. Design of indexes for optimum performance
- VI. Event and snapshot Monitors**
 - A. How to use event monitors
 - B. Options for creating and using event monitors
- VII. Multi-Dimensional Clustering (MDC)**
 - A. The terminology of multi-dimensional data
 - B. Know why and how to create a table that is clustered in multiple dimensions.
 - C. Create generated columns to be used for dimension columns.
 - D. Understand the concept of monotonicity associated with generated dimension columns.
- VIII. Advanced table options**
 - A. Using table range partitions
 - B. Table row compression
 - C. Advanced options for REORG, RUNSTATS utilities
 - D. Supporting identity columns in utilities
- IX. Configuring for self tuning of DB2**
 - A. Automatic maintenance with utilities
 - B. Self tuning memory
 - C. Database global memory
 - D. Automatic tuning
 - E. DB2 Operational Model
 - F. Review the architectural concepts for DB2
 - G. Look at the behavior of DB2, with respect to I/O operations
 - H. Identify important cache and monitoring elements
- X. The Health Monitor and Activity Monitor**
 - A. Using commands and the Health Center to monitor DB2 "health"
 - B. Monitor DB2 memory resource usage with the Activity Monitor
 - C. Various tools, like DB2PD will be used to monitor memory utilization
 - D. Analyzing performance with snapshot functions and administrative views
- XI. DB2 Configuration Parameters**
 - A. Review important configuration parameters for the a database and the database manager
 - B. Introduce auto-configuration
 - C. Identify the concepts of the Design Advisor