Advanced SQL for DB2 UDB for LUW

Advanced SQL for DB2 UDB for LUW (Linux, UNIX, Windows) is designed for the SQL professional who wants to get more "bang for the buck" out of DB2, especially DB2 Universal Database. A common complaint is that so much money is sunk into relational database management systems, only to have them end up as glorified flat-file processors due to inefficient SQL and misguided application design.

Why select all the columns and all the rows only to have your application program (or third-party tool) do the filtering? Why select detail-level rows when what you want to see is summarized results? And what about all the new advanced features added to SQL on the DB2 Universal Database platforms?

This course presents advanced SQL syntax in an easy-to-understand format, demystifying even the most complex of the new features. It enables the attendee to return to work much better prepared to avail him- or herself of the full power of the DB2 engine.

Audience

Application programmers, database administrators, sophisticated end users, and anyone who needs to know more about the intricacies of SQL.

Prerequisites

Before taking this course, you should be able to code basic *to intermediate* SQL statements and have real-world experience coding such statements.

You can develop these skills by attending either *Introduction to DB2 and SQL* or *DB2 and SQL Fundamentals*.

Proficiency with Windows, Notepad (or another Windows-based text editor), the mouse, and the PC in general is also assumed.

Outline/Schedule

Day 1

Chapter 1—Review of SQL Fundamentals

Workshop 1—Familiarization and Review

Chapter 2—Subqueries

Workshop 2—Subqueries

Chapter 3—Two-dimensional Matrices

Workshop 3—VALUES, Table Expressions, and Set Operators

Day 2

Chapter 4—Expressions and Scalar Functions

Workshop 4—Expressions and Scalar Functions

Chapter 5—Inner and Outer Joins

Workshop 5—Inner and Outer Joins

Chapter 6—Aggregation

Workshop 6—Grouping Sets and Super Groups

Day 3

Chapter 7—Recursive SQL

Workshop 7—Recursive SQL

Chapter 8—Object-Relational Features

Workshop 8—User-defined Database Objects

Chapter 9—Performance Considerations

Workshop 9—Visual Explain

Chapter 10—Review and Summation

Objectives

Upon completing this course, you should be able to:

- Be fully conversant in relational concepts and terminology, and code simple to complex SELECT, INSERT, UPDATE, and DELETE statements
- Differentiate between the syntax elements of *subselect*, *fullselect*, and *select-statement*, and the terms *query* and *subquery*, and code both correlated and noncorrelated subqueries
- Identify various types of two-dimensional inputs to a query, including the two-dimensional output of a prior query, and combine and contrast multiple two-dimensional results with the UNION, INTERSECT, and EXCEPT set operators
- Use advanced expressions and scalar functions, as well as concatenations and calculations, including sophisticated date-time manipulation and arithmetic
- Differentiate between and code inner joins and left, right, and full outer joins, as well as Cartesian products
- Perform aggregate processing using advanced column functions, grouping sets, and the super groups ROLLUP and CUBE
- Code recursive SQL using common table expressions or views, and identify and avoid situations that cause infinite loops
- Code SQL that utilizes the vast array of object-relational features available in DB2 Universal Database, including data-integrity constraints, user-defined distinct types (UDT's), user-defined functions (UDF's), triggers, stored procedures, and compound SQL
- Adhere to some simple guidelines to ensure that your SQL—whether extremely basic or impossibly complex—performs as well as possible