

Training: IBM DB2 SQL Workshop for Experienced Users



TRAINING GOALS:

This course teaches you how to make use of advanced SQL techniques to access DB2 databases in different environments. This course is appropriate for customers working in all DB2 environments, specifically for z/OS, Linux, UNIX, and Windows.

- Discuss basic relational database concepts
- Use some of the OLAP features of DB2, such as GROUPING and RANKING functions
- Create tables, views and indexes
- Use referential integrity, check constraints and triggers
- Use outer joins, and join tables to themselves
- Use CASE expressions, and the CAST function
- Identify the impact of Summary Tables, Materialized Query Tables, and temporary tables
- Use complex subqueries
- Use a greater number of scalar SQL functions
- Use advanced SQL constructs, such as recursive SQL and table expressions
- Define User-Defined Distinct Types and User-Defined Functions
- Avoid several of the most common causes for poorly-performing SQL

This intermediate course is for experienced SQL end users, application programmers, database administrators, and user support staff who need more advanced knowledge of SQL.

CONSPECT:

Introduction

- Identify the purpose of the clauses in the SELECT statement
- Describe the key differences among the IBM DB2 platforms
- Describe and use some of the OLAP features of DB2, such as GROUPING functions like CUBE and ROLLUP, and the RANK, DENSE_RANK and ROW_NUMBER functions

Create Objects

- Code statements to: Create tables and views, Alter tables, Create indexes, Implement

referential integrity (RI), and Define triggers and check constraints

- Identify impacts and advantages of referential integrity, including impacts of delete rules
- Identify considerations when using triggers and check constraints
- Define and make use of INSTEAD OF triggers

Join

- Retrieve data from more than one table via inner and outer joins
- Use outer joins (LEFT, RIGHT, FULL)
- Use ANTI JOINS
- Join a table to itself
- Use UNION and UNION ALL
- Use EXCEPT and INTERCEPT

CASE, CAST, Summary Tables, and Materialized Query Tables

- Identify when CASE expressions can be used
- Code CASE expressions in SELECT list and in the WHERE clause
- Identify when CAST specifications can be used
- Identify the advantages of using Summary (Materialized Query) Tables and Temporary tables
- Identify the advantages of using Materialized Query Tables (MQTs)
 - Identify when and how to use Temporary tables

Using Subqueries

- Code subqueries using the ALL, ANY/SOME, and EXISTS keywords
- Code correlated subqueries
- Choose the proper type of subquery to use in each case

Scalar Functions

- Extend your knowledge of scalar functions which: Manipulate arithmetic data, Manipulate date values, and Manipulate character data
- Examples of scalar functions that are addressed in this course:
 - SUBSTR
 - POSSTR
 - COALESCE/VALUE
 - DECIMAL
 - ROUND
 - DIGITS
 - CHAR

- DATE/TIME

Table Expressions and Recursive SQL

- Identify reasons for using table expressions and recursive SQL
- Use nested and common table expressions
- Identify the difference between views and table expressions
- Code recursive SQL
- Control the depth of recursion when coding recursive SQL

UDTs/UDFs and Performance

- Describe the concepts behind User-Defined Types, User-Defined Functions and Stored Procedures
- Predict when queries will use indexes to get better performance
- Identify concepts of predicate processing
- State introductory concepts about index structure
- State general best practices advice

REQUIREMENTS:

You should have experience with:

- coding **and** executing basic SQL statements.

These skills can be developed by attending

- *SQL Workshop (CE120)*, **or** equivalent experience.

Difficulty level

