IBM Db2 12: SQL Workshop for Experienced Users

Duration: 3 Days Course Code: CLA32G Delivery Method: Virtual Learning

Overview:

This course covers fundamental and advanced SQL techniques for efficient data management, optimization, and analytics in Db2 12. Participants will learn how to navigate the Db2 system catalog and explore query performance optimization techniques, including indexing strategies, predicate pushdown, and query access plans to improve efficiency.
The course also introduces advanced SQL functions, such as LISTAGG, PERCENTILE_CONT, and PERCENTILE_DISC. Participants will learn how to create and manage database objects and gain expertise in utilizing advanced join techniques such as inner and outer joins, hash joins, lateral joins, and set operators like UNION, INTERSECT, and EXCEPT to merge and analyze data from multiple tables.
The course also covers the use of subqueries and query optimization techniques, including correlated subqueries, EXISTS/NOT EXISTS, and quantified predicates to enhance data retrieval efficiency.
Through hands-on exercises and real-world scenarios, they will develop the skills necessary to handle large datasets efficiently and implement best practices for high-performance database management.

Virtual Learning

This interactive training can be taken from any location, your office or home and is delivered by a trainer. This training does not have any delegates in the class with the instructor, since all delegates are virtually connected. Virtual delegates do not travel to this course, Global Knowledge will send you all the information needed before the start of the course and you can test the logins.

Target Audience:

This course is ideal for database administrators, SQL developers, and data analysts who have prior experience with SQL and wish to enhance their Db2 v12 expertise.
This course is appropriate for customers working in Db2 environments for Linux, UNIX, and Windows.

Objectives:

- After completing this course, learners should be able to:
- Demonstrate the use of OLAP ranking functions (RANK, DENSE_RANK, ROW_NUMBER) to analyze and order query results
- Optimize queries with expression-based indexes for computed values
- Utilize set operators like UNION, EXCEPT, and INTERSECT for advanced data merging
- Leverage summary tables and MQTs to precompute and store aggregated data
- Apply EXISTS, NOT EXISTS IN, IN, and NOT IN predicates for conditional data retrieval
- Perform numeric and mathematical operations, including logarithmic, exponential, and trigonometric calculations
- Implement recursive SQL to process hierarchical relationships
- Leverage ARRAY data types to pass structured lists of values within SQL queries and procedures
- Analyze and optimize SQL execution plans with EXPLAIN and VISUAL EXPLAIN

Prerequisites:

Before taking this course, participants should:

- Be able to discuss basic Relational Database concepts
- Be able to code basic SQL statements

OR

- Have completed the SQL Workshop course (CLA22G or 2LA22G)

Content:

Unit 1: SQL Introduction

- Describe the structure and purpose of Db2 v12, including its system catalog, database organization, and key SQL components
- Explain how Db2 catalogs and system tables store metadata, and how they can be queried for database object information
- Demonstrate the use of OLAP ranking functions (RANK, DENSE_RANK, ROW_NUMBER) to analyze and order query results

Unit 2: Create Objects

- Create and modify tables, columns, and constraints
- Use generated columns to automate calculations and reduce dependency on triggers
- Optimize queries with expression-based indexes for computed values
- Implement referential integrity, check constraints, and triggers for enforcing business rules
- Leverage views and materialized query tables to simplify data access and reporting

Unit 3: JOIN

- Use inner and outer joins to retrieve related data across tables
- Optimize query performance with hash joins for large datasets
- Leverage lateral joins for more flexible row-by-row processing
- Perform self-joins to analyze hierarchical or recursive relationships
- Utilize set operators like UNION, EXCEPT, and INTERSECT for advanced data merging

Unit 4: CASE, CAST, Summary (Materialized Query) Tables, and Temporary Tables

- Use CASE expressions to apply conditional logic within queries
- Perform data type conversions using CAST
- Leverage summary tables and MQTs to precompute and store aggregated data
- Implement temporary tables and SGTTs to handle intermediate results

Unit 5: Using Subqueries

- Use basic subqueries to filter and retrieve data dynamically
- Leverage correlated subqueries for row-dependent queries
- Apply EXISTS, NOT EXISTS, IN, and NOT IN predicates for conditional data retrieval
- Utilize LATERAL subqueries to execute row-by-row subqueries for sequential logic processing
- Optimize query performance by structuring subqueries efficiently

Unit 6: Scalar Functions

- Use string functions to extract, modify, and format text data
- Perform numeric and mathematical operations, including logarithmic, exponential, and trigonometric calculations
- Apply LISTAGG to concatenate multiple row values into a single string
- Work with data and time functions to extract, manipulate, and format timestamps

Unit 7: Table Expressions and Recursive SQL

- Use nested table expressions to create temporary result sets for complex queries
- Leverage common table expressions (CTEs) for modular and reusable SQL structures
- Implement recursive SQL to process hierarchical relationships
- Apply performance optimization techniques to improve efficiency in recursive queries

Unit 8: An Introduction to UDTs / UDFs and Stored Procedures

- Define User-Defined Distinct Types (UDTs) to create custom data types for specialized business needs
- Implement User-Defined Functions (UDFs)
- Utilize stored procedures to encapsulate complex database logic for performance efficiency
- Leverage ARRAY data types to pass structured lists of values within SQL queries and procedures

Unit 9: SQL and Db2 Performance

- Use indexes, clustering, and predicate pushdown
- Analyze and optimize SQL execution plans with EXPLAIN and VISUAL EXPLAIN
- Improve filtering and sorting performance using predicate processing techniques
- Leverage RUNSTATS enhancements to ensure up-to-date statistics for query optimization

Additional Information:

Official course book provided to participants.

Further Information:

For More information, or to book your course, please call us on +1 (800) 268 7737

ilt.customerservice@globalknowledge.com

www.globalknowledge.com/en-ca/

Global Knowledge Network (Canada) Inc., 150 Knowledge Park Drive, Suite 130, Fredericton, NB, E3C 0N9