

## Educational Seminar Details

---

### Instructor Information:

Sheryl Larsen  
President, Sheryl M. Larsen, Inc.  
409 Hawthorne Avenue  
Elmhurst, IL 60126  
(630) 399-3330  
(630) 834-0069 fax  
[SherylMLarsen@cs.com](mailto:SherylMLarsen@cs.com)  
[www.smlsql.com](http://www.smlsql.com)

### Instructor Biography:

Sheryl Larsen is an internationally recognized researcher, consultant and lecturer, specializing in DB2 and is known for her extensive expertise in SQL. Sheryl has over 18 years experience in DB2, has published many articles, several popular DB2 posters, and co-authored a book, DB2 Answers, Osborne-McGraw-Hill, 1999. She was voted into the IDUG "Speaker Hall of Fame" in 2001 and was the Executive Editor of the IDUG Solutions Journal magazine 1997-2000. Currently, she is President of the Midwest Database Users Group ([www.mwdug.org](http://www.mwdug.org)), a member of IBM's DB2 Gold Consultants program, and owns Sheryl M. Larsen, Inc. ([www.smlsql.com](http://www.smlsql.com)), a firm specializing in Advanced SQL Consulting and Education.

**Seminar Title:** DB2 Advanced & Complex SQL

**What audience level is this seminar geared for?**  Beginner  Intermediate  x Advanced

### Goals of the Seminar:

1. Learn where to use what SQL so that you can efficiently pre-process your data and react to the fast changing business world.
2. Be able to implement the sample queries provided into your business environment
3. Get updated on the latest release SQL enhancements to better prepare your applications for the future

### Abstract:

Performance is measured by response time, throughput, peak response time, hits, and sessions per second. SQL coding techniques directly impact performance. Prepare yourself for developing high performing DB2 UDB applications on any platform. Leave the presentation with an understanding of the many sample queries provided and the knowledge of how to leverage their technique at your shop. Sample queries include complex relational division, recursive query technique, global temporary tables, dynamic SQL, CASE expressions, row expressions, limited fetch, scrollable cursors, and UNION everywhere. DB2 Version 8 for OS/390 SQL enhancements, all 20, will also be discussed in detail.

**Where and when this class was previously presented?**



## Seminar Outline

- |  |   |
|--|---|
| <p>I. SQL Performance</p> <ol style="list-style-type: none"> <li>a. DB2 Engine Components</li> <li>b. Predicate Processing Intelligence</li> <li>c. Optimizer Overview</li> <li>d. Access Path Overview</li> </ol> <p>II. SQL Environmentals</p> <ol style="list-style-type: none"> <li>a. Good SQL Coding Standards</li> <li>b. Clear SQL Documentation Techniques</li> <li>c. Building Complex SELECTs</li> </ol> <p>III. Complex Relational Division</p> <ol style="list-style-type: none"> <li>a. Joins</li> <li>b. Subselects</li> <li>c. Table Expressions</li> <li>d. Outer Joins.</li> </ol> <p>IV. Recursive queries</p> <ol style="list-style-type: none"> <li>a. Basic syntax</li> <li>b. How does Recursion Start</li> <li>c. How does Recursion Stop</li> <li>d. Emulating Recursion on OS/390</li> </ol> <p>V. Global Temporary Tables</p> <ol style="list-style-type: none"> <li>a. Created</li> <li>b. Declared</li> <li>c. Practical Uses</li> </ol> <p>VI. Dynamic SQL</p> <ol style="list-style-type: none"> <li>a. Basic Syntax</li> <li>b. Emulating</li> <li>c. Complex Examples</li> </ol> <p>VII. CASE</p> <ol style="list-style-type: none"> <li>a. Basic Syntax</li> <li>b. The True Power</li> <li>c. Complex Examples</li> </ol> | <p>VIII. SQL Functions (the practical ones)</p> <ol style="list-style-type: none"> <li>a. Built-in</li> <li>b. Scalar</li> <li>c. The True Cost</li> <li>d. Complex Examples</li> </ol> <p>IX. DB2 for OS/390 and z/OS Version 7 SQL Enhancements</p> <ol style="list-style-type: none"> <li>a. Limited Fetch</li> <li>b. Scrollable Cursors</li> <li>c. UNION Everywhere</li> <li>d. Row Expressions</li> <li>e. MIN/MAX Single Index Support</li> <li>f. Self Referencing Updates with Subselects</li> <li>g. Sort Avoidance for ORDER BY</li> <li>h. Bind Improvements</li> </ol> <p>X. DB2 V8 SQL Enhancements</p> <ol style="list-style-type: none"> <li>a. Stage1 unlike data types</li> <li>b. 2M Statement Length</li> <li>c. Multit-row INSERT</li> <li>d. Multit-row FETCH</li> <li>e. Dynamic Scrollable Cursors</li> <li>f. GROUP BY Expression</li> <li>g. Materialized Query Tables</li> <li>h. INSERT with SELECT</li> <li>i. Parallel Sort</li> <li>j. Multiple DISTINCT</li> <li>k. Sequences</li> <li>l. Star Join Sparse Index</li> <li>m. Multiple CCSIDs per statement</li> <li>n. VOLATILE Table Support</li> <li>o. Enhanced UNICODE</li> <li>p. GET DIAGNOSTICS</li> <li>q. CURRENT PACKAGE PATH</li> <li>r. Qualified Column names</li> <li>s. Common Table Expressions</li> <li>t. Recursive SQL</li> </ol> |
|--|---|

**Seminar Title: Tuning DB2 SQL for Performance**

**What audience level is this seminar geared for?**  Beginner  Intermediate  Advanced

**Goals of the Seminar:**

1. Learn about *all* access paths that can be chosen by the DB2 UDB for OS/390 optimizer
2. Learn when and how to alter optimizer chosen access paths
3. Learn how to optimize your index design

**Abstract:**

Companies now recognize that IT projects must deliver measurable value. Performance is measured by response time, throughput, peak response time, hits, and sessions per second. Index design, access path selection, and table join sequence, directly impact performance. Come learn how to take control of your performance! Once proactive, the next concern will be where to shift the savings.

This is a detailed review of all access paths that can be chosen by the DB2 UDB for OS/390 optimizer. Access paths will be shown in animated detail while discussing when the optimizer prefers each access path and how to encourage or discourage each. This seminar will also take a look at practical methods for altering optimizer chosen access paths. IBM warranted tuning methods such as OPTIMIZE FOR 1 ROW, no ops, REOPT(VARS), as well as creative methods, such as fake filtering and table expressions, will be demonstrated. The how to tune SQL is easy, it is the when that is the challenge. Index design also directly impacts performance, however, lack of staff and financial resources usually puts proactive index design on the bottom of the IT To-Do-List. Leave the seminar with a proven technique for optimizing index design.

**Where and when this class was previously presented?**

Date: 2000 - 2003

Location: Worldwide

**Seminar Outline****•Tuning Queries**

–When? Why? How?

**•Learn all the possible access paths**

–Variations of index access

–Variations of table access

–Variations of join methods

**•Check the access paths of each query block**

–Using EXPLAIN

–Using Visual EXPLAIN

**•Learning Explain**

–IBM's PLAN\_TABLE definition

–My PLAN\_TABLE definition

–PLAN\_TABLE Access Path Identifiers

**•Index Accesses****•Table Accesses****•Partition Accesses****•Hidden Processes****•Apply data knowledge and program knowledge to predict response time**

– Estimating Example 1 – Correlated Subquery

– Estimating Example 2 – List Prefetch

– Estimating Example 3 – Join Sequence

•If response time needs improvement, Tune using a proven method

-Tuning Example 1 – OPTIMIZE FOR n ROWS/FETCH FIRST n ROWS ONLY

-Tuning Example 2,3,4 – No Operations

-Tuning Example 5, 6 – Fake Filtering

-Tuning Example 7 – Table Expressions

-Tuning Example 8 – Index Design

•Summary