

# Examples of Db2 Mistakes for Others to Avoid

Sheryl M. Larsen

**Kyndryl**

Kyndryl, zCloud Sales



1

**Notes:** This presentation divulges discoveries and recommendations from various SQL performance review assignments. See if you have similar SQL performance issues and get the instructions on how to fix them. Issues discussed include non-optimal index design, access paths gone wild, Stage 2 predicates, misuse of SQL, excessive sorting, delayed filtering, lack of implementation of powerful new SQL features.

#### Skills Taught:

- Be able to determine *where* to use *what* SQL features, for example, can identify when to use joins versus subqueries and vice versa.



# First SQL Class DB2 V1 beta



```
SELECT      O.ORDERID, C.CUSTOMERID,
            B.BILL, SUM(B.AMOUNT) AS TOTAL
FROM        ORDER O, CUSTOMER C, BILL B
WHERE       B.DATE > '01-01-2017'
            AND O.ORDERID = C.ORDERID
            AND C.CUSTOMERID = B.CUSTOMERID
GROUP BY    O.ORDERID, C.CUSTOMERID, B.BILL
HAVING      TOTAL > 100000
ORDER BY    TOTAL DESC
```



**Notes:**

# 1984

- No internet
- No cell phone
- No laptop
- No ear buds
- No email
- No .....



- rldicl (Rotate Left Double Word Immediate then Clear Left) instruction
- rlmi (Rotate Left Then Mask Insert) instruction
- rrib (Rotate Right and Insert Bit) instruction
- sld (Shift Left Double Word) instruction
- sle (Shift Left Extended) instruction
- sleg (Shift Left Extended with MQ) instruction
- sliq (Shift Left Immediate with MQ) instruction
- slliq (Shift Left Long Immediate with MQ) instruction



**Notes:**

# Fast Forward



35  
Y  
E  
A  
R  
S

1984 **Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, ORDER BY,** Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join, Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, SQL PL in routines, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, MERGE not logged tables, OmniFind, spatial, range partitions, data compression, DECFLOAT, optimistic locking, ROLE, TRUNCATE, index & XML compression, created temps, inline LOB, administrative privileges, implicit cast, increased timestamp precision, currently committed, moving sum & average, index include columns, row and column access controls, time travel query, GROUPING SETS, ROLLUP, CUBE, global variables, Text Search functions, accelerated tables, DROP COLUMN, array data type, XML enhancements, moving SUM/AVG, Array variables, ARRAY\_EXISTS, COUNTBIG, 2019  
2020  
2021 SELECT INTO statements with UNION or UNION ALL allowed  
SELECT INTO statements with UNION or UNION ALL disallowed, OFFSET, 11 global variables, LISTAGG + 33 more BIFs, LIMIT fetch row count1



**Notes:**

# Case Studies

## Lessons Learned From SQL Performance

### Reviews:

- Lack of Use of Powerful New and Old SQL Features
- Misuse of SQL
- Using Wrong Type of Temp Table
- Non-Optimal Index Design
- Seriously Delayed Filtering
- Access Paths Gone Wild
- Excessive Sorting
- Performance Structures Missing
- Stage 2 Predicates in Frozen SQL



**Notes:** The **Lessons Learned from SQL Performance Review** presentation divulges discoveries and recommendations from various SQL performance review assignments. Come see if you have similar SQL performance issues and get the instructions on how to fix them. Issues discussed include non-optimal index design, access paths gone wild, Stage 2 predicates, misuse of SQL, excessive sorting, delayed filtering, lack of implementation of powerful new SQL features.

# Client SQL Report Card

Static SQL	Dynamic SQL
V1 – 251 GROUP BY	V1 – 8 GROUP BY
V1 – 26 Subqueries	V1 – Zero Subqueries
V2 – 380 OPTIMIZE FOR 1 ROW	V2 – Zero OPTIMIZE FOR 1 ROW
V4 – ? Created Global Temp Tables	V4 – Zero Created Global Temp Tables
V4 – 25 Table Expressions	V4 – Zero Table Expressions
V4 – 520 WITH UR	V4 – 12 WITH UR
V5 – 951 CASE Expressions	V5 – Zero CASE Expressions
V6 – Zero Declared Temp Tables	V6 – Zero Declared Temp Tables
V7 – 406 FETCH FIRST n ROWS	V8 – 2 FETCH FIRST n ROWS
V8 – 1 SELECT INTO with ORDER BY	V8 – Zero SELECT INTO with ORDER BY
V8 – 545 Common Table Expressions	V8 – Zero Common Table Expressions
V8 – 2 Multi-row Fetch	V8 – Zero Multi-row Fetch
V8 – ? Materialized Query Tables	V8 – Zero Materialized Query Tables
V9 – 69 EXCEPT/INTERCEPT	V8 – Zero EXCEPT/INTERCEPT
V9 – 61 MERGE/TRUNCATE	V9 – Zero MERGE/TRUNCATE
V9 – Zero RANK/DENSE_RANK/ROW_NUMBER	V9 – Zero RANK/DENSE_RANK/ROW_NUMBER

© Sheryl M. Larsen, Inc. 2000-2023 7

**Notes:** Some use of SQL technology but still used CURSORS for fetching 1 ROW with ORDER BY.

Name: \_\_\_\_\_

# SQL Coding Skill Self Assessment

Before: \_\_\_\_\_

*Basic  
SQL Skill  
Range*

*Intermediate  
SQL Skill  
Range*

*Advanced  
SQL Skill  
Range*

After: \_\_\_\_\_

Level	Assessment = YOU CAN FULLY UNDERSTAND THE FEATURE AND PROPER USE OF:
0	You think SQL is a new energy drink
1	Simple SELECT statements, WITH clause, ORDER BY
2	WHERE clauses, BETWEEN, LIKE, IN(list), =, >=, >, <, <=, <>, NOT IN(list), NOT LIKE, NOT BETWEEN
3	Table joins (inner, outer, full), UNION, UNION ALL, CONCAT, static CURSORS, FOR UPDATE OF, ROW_NUMBER
4	noncorrelated and correlated subqueries, EXISTS, NOT EXISTS, FETCH FIRST x ROWS ONLY, OPTIMIZE FOR x ROWS
5	Indexable, Stage1 and Stage 2 predicate evaluation, multirow FETCH/INSERT, GET DIAGNOSTICS, Scalar full SELECT,
6	Table expressions/common table expressions, GROUP BY, HAVING, IS NOT DISTINCT FROM, EXCEPT/INTERCEPT
7	CASE expressions, Global Temporary Table (GTT), Declared Temporary Table (DTT), Dynamic Scrollable cursors, SEQUENCE columns
8	Queries involving > 10 tables, INSERT within SELECT, Star Schema, Snow Flake, GROUP BY expression, IDENTITY columns
9	MQT (Materialized Query Tables), Recursive SQL, UNION in Views, Native SQL Stored Procedures, > 20 SQL Functions, DENSE_RANK, RANK
10	Codes effective and efficient SQL applying performance rules and knows when to use each appropriately



**Notes:**

# Dynamic Queries

Had no Sheriff



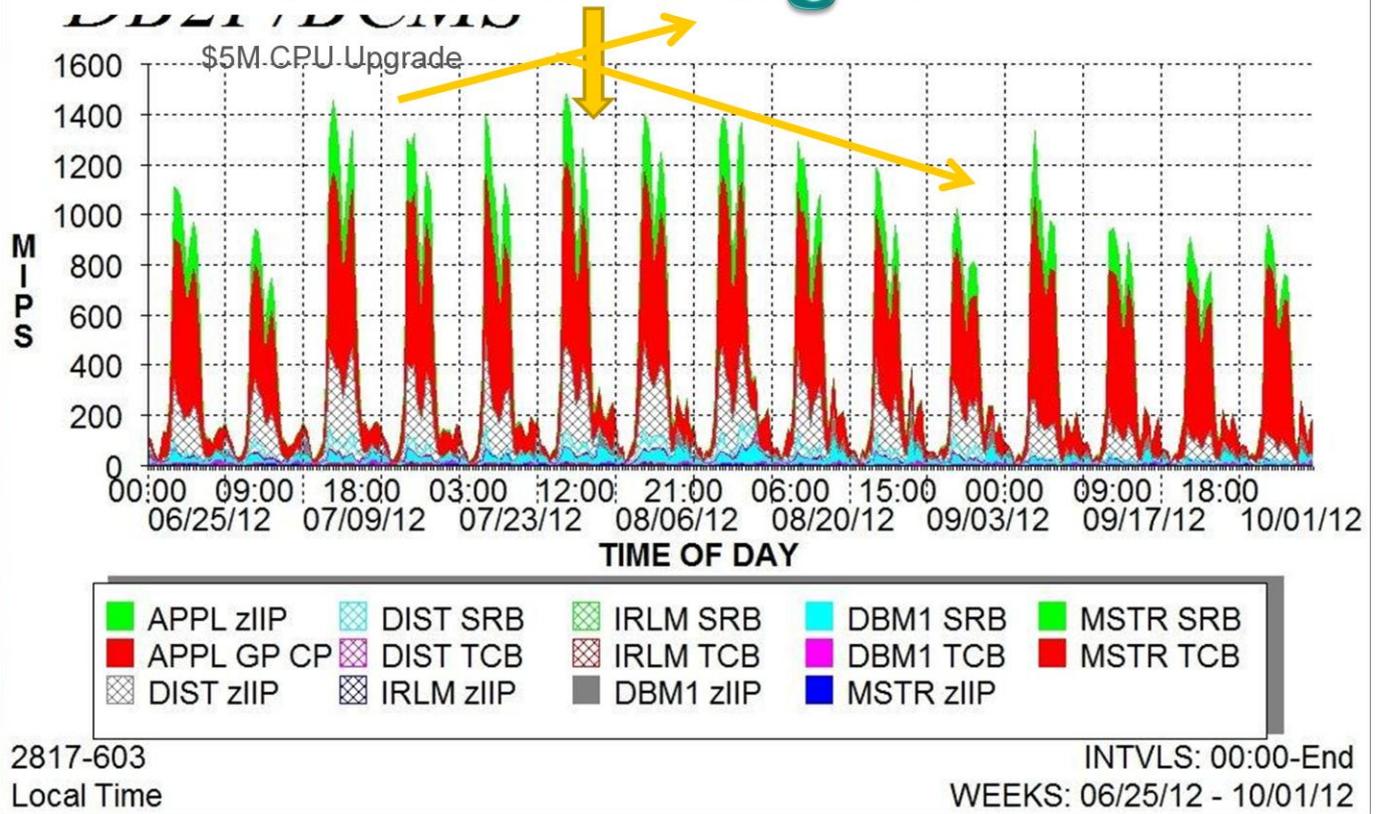
Is Db2 Up?



**Notes:**

# DB2 CPU MIPS

## First Changes



### Notes:

# Visual Plan Graphs - Bad

The screenshot displays the IBM Data Studio interface with the following components:

- Configuration:** Connection: PRODDB2P, Run method: JDBC.
- SQL Query:**

```

COALESCE ( SEX,
TX_SUB_MCR_D_EL,
TX_SUB_LAST_CHN,
TX_SUB_EMPTY_ST,
TSEQ.EMPL_OCC_D
FROM CIMP_TB_TXN_SUB
ON TX_SUB_TXN_I
TX_MSTR_SUB_ID
MRTL_MRTL_STAT
SEX ON SEX_SEX
CIMP_TB_TXN_S
AND TSEQ.TXN
CIMP_TB_TXN_F
FSD.TXN_TS =
WHERE ( TX_MSTR.TXN
TX_SUB.SUB_ID

```
- Visual Plan Graph:** A complex tree structure of nodes. Key nodes include:
  - Root: @@HJOIN 0
  - Intermediate: @@FETCH 1, @@XSCAN 1, @@TB\_MRTL\_STAT\_CD 3, @@XSEXD01 3
  - Leaf nodes: @@TBSCAN 1,0204, @@XSCAN 1, @@TB\_BAD\_ADDR 54620, @@XBADDR2 53437, @@QB5, @@QB7, @@QB8, @@QB9, @@QB10, @@QB11
- Attributes Table:**

Name	Output Cardinality	Cumulative Total Cost	Cumulative IO Cost	Cumulative CPU Cost	Runtime Parent Query Bloc	Do at Open	Times
@@HJOIN 0							
@@FETCH 1							
@@XSCAN 1							
@@TB_MRTL_STAT_CD 3							
@@XSEXD01 3							
@@TBSCAN 1,0204							
@@XSCAN 1							
@@TB_BAD_ADDR 54620							
@@XBADDR2 53437							
@@QB5							
@@QB7							
@@QB8							
@@QB9							
@@QB10							
@@QB11							

STMT_ID	CPU	EXECS	ELAPSED	GETPAGES	EXAMINED ROWS	PROCESSED	STAT_SORT	STAT_INDX	STAT_RSCN
1022787	4705.7	5258	4722.46	729,475,051	298,690,230	2629	0	283543230	5606



© Sheryl M. Larsen, Inc. 2000-2023

Notes:

# Removed Bad Scalar Subquery

Query now just returns a 'Y'  
First rewrite implemented  
September 19th

```
SELECT 'Y'  
FROM ABCP.TB_BAD_ADDR ADDR  
  INNER JOIN  
    ABCP.TB_GRP_ELEC_ONLY XGRP  
  ON XGRP.GRP_ID = ADDR.GRP_ID  
  INNER JOIN ABCP.TB_SUB_CONT SUB  
  ON SUB.SUB_ID = ADDR.SUB_ID  
    AND SUB.GRP_ID = ADDR.GRP_ID  
WHERE ADDR.ADDR_RSLV_FL = 'N'  
  AND XGRP.EMPR_ID = 17469  
  AND ( SUB.SUB_CONT_TERM_DT > (CURRENT  
DATE - 90 DAYS) )  
FETCH FIRST 1 ROW ONLY  
FOR READ ONLY
```

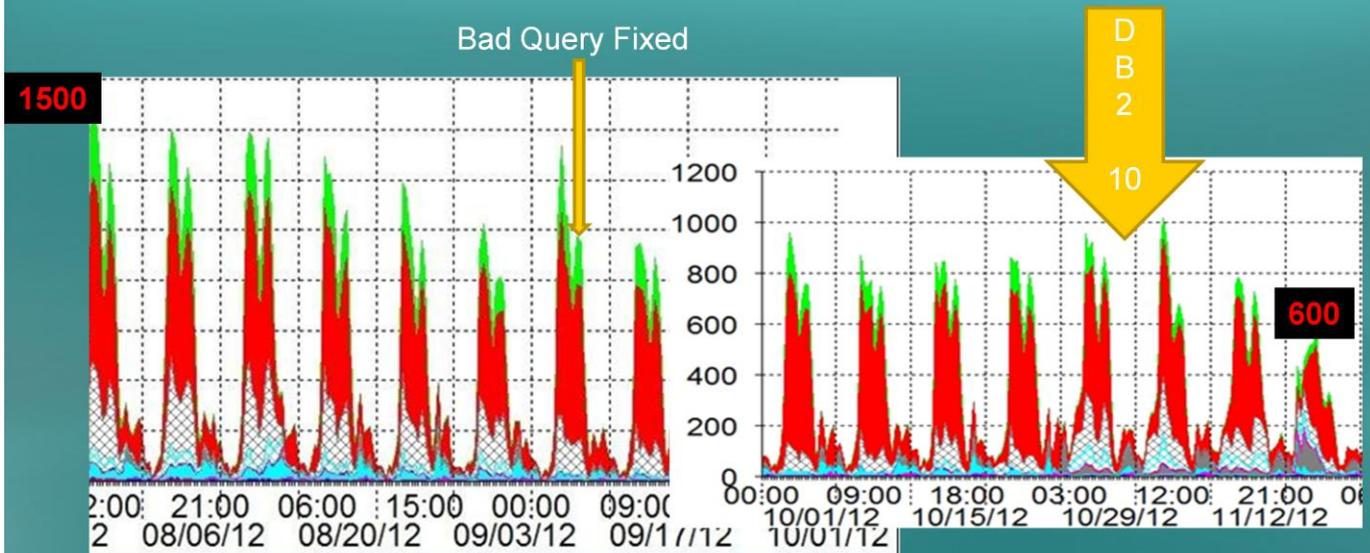
USE COUNT	-> 2	SQL	-> 4
TIMEPCT	-> .00%	CPUPCT	-> .00%
INDB2 TIME	-> 00:00.033008	INDB2_CPU	-> 00:00.012755
<b>GETPAGE</b>	<b>-&gt; 642</b>	GETPFAIL	-> 0
SYNCREAD	-> 2	SPFETCH	-> 0
LPFETCH	-> 0	DYNPFETCH	-> 85
PFPAGES	-> 11	PAGEUPDT	-> 0

**Notes:** Performance rule violations usually result in increased CPU or I/O, time to fix the mistake, and ultimately, a cost to the business unit.

What will the cost be? Depends on the mistake and the frequency of the mistake.

The following presentation will show real case studies of the actual cost of the mistakes.

# DB2 CPU Reduction



## 900 MIP Reduction



**Notes:**

# Misuse of SQL Costs \$\$\$\$\$\$\$\$\$

Misuse has a cost for:

JOINs

Subqueries

LEFT JOINs

GLOBAL TEMPORARY TABLEs



**Notes:**

# Joins over Subqueries

- ◆ When detail row information is required
  - » CLASS\_ID, TIME\_ID, PHONE\_NO
  - » What Students Have the Same Classes as any given student?



Point to  
Different  
Rows

```
SELECT SSS.SID, SSS.CLASS_ID, SSS.TIME_ID
FROM   SMLU_STUDENT_SCHED SS
       ,SMLU_STUDENT_SCHED SSS
WHERE  SSS.CLASS_ID = SS.CLASS_ID
       AND SSS.TIME_ID = SS.TIME_ID
       AND SSS.SID <> SS.SID
       AND SSS.SID NOT IN(list challenged student IDs)
       AND SS.SID = :challenged-sid
```

Self Join to get same  
class/time rows

Homework  
Challenged



© Sheryl M. Larsen, Inc. 2000-2023

15

**Notes:** Joins are very good at giving lots of detail from either table.

# Subqueries over Joins

- ◆ When detail row information is not required
  - » Unique SIDs only
  - » What Students Have **at least one** of the same Classes as a given student?

```
SELECT SS.SID
FROM SMLU_STUDENT_SCHED SS
WHERE SS.SID NOT IN(list challenged student IDs)
AND EXISTS
  (SELECT 'TRUE OR FALSE'
   FROM SMLU_STUDENT_SCHED SSS
   WHERE SSS.CLASS_ID = SS.CLASS_ID
        AND SSS.TIME_ID = SS.TIME_ID
        AND SSS.SID <> SS.SID
        AND SSS.SID = :challenged-sid)
```

**Homework  
Challenged**



**Notes:** Correlated subqueries are good at bypassing detail.

# Disallow DISTINCT Joins!

```
SELECT      DISTINCT SSS.SID
FROM        SMLU_STUDENT_SCHED SS
           ,SMLU_STUDENT_SCHED SSS
WHERE       SSS.CLASS_ID = SS.CLASS_ID
           AND SSS.TIME_ID = SS.TIME_ID
           AND SSS.SID <> SS.SID
           AND SSS.SID NOT IN(list challenged student IDs)
           AND SS.SID = :challenged-sid
```



© Sheryl M. Larsen, Inc. 2000-2023

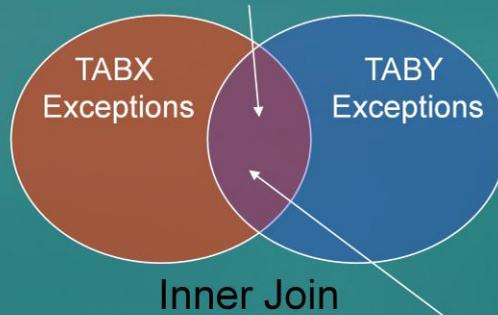
17

**Notes:** DISTINCT should never be used with an INNER JOIN. The INNER JOIN produces detail and the DISTINCT removes the detail.

# Inner Join over Left Join

```
SELECT Columns  
FROM TABY Y  
LEFT JOIN TABX X  
ON X.COL1 = Y.COL1  
WHERE Y.COL1 IS NOT NULL
```

**Slower**



**Faster**

Does not apply to LUW

```
SELECT Columns  
FROM TABX X, TABY Y  
WHERE X.COL1 = Y.COL1
```



**Notes:** This INNER JOIN is faster than this LEFT JOIN when there are many rows that do not join. This is because LEFT JOINS have to calculate exceptions to the JOIN.

# Using Wrong Type of Temp Table

Created versus Declared



**Notes:**

# CRM Application Package

(SELECT Lots of Columns  
FROM



INNER  
JOIN



INNER  
JOIN



(SELECT  
FROM



LEFT JOIN

(SELECT  
FROM



LEFT JOIN

(SELECT DISTINCT  
FROM



GROUP BY

- 
- Rapid Application Development (outsourced) treated all transactions equal
  - No internal performance review for high volume extracts
  - Application consumed 25% of new z/9 on a daily basis
  - Most expensive program running had multi-step Temp Table Processing



**Notes:** A problem application consumed computer and people resources when scaling was attempted.

# Business Deadline Driven

- Lack of SQL skill of developers hindering performance
  - » Misuse of SQL features
    - Used Created Temp Tables, CTT instead of Declared Temp Table, DTT which allows indexes

SELECT  
FROM

CTT

SELECT  
FROM

DTT

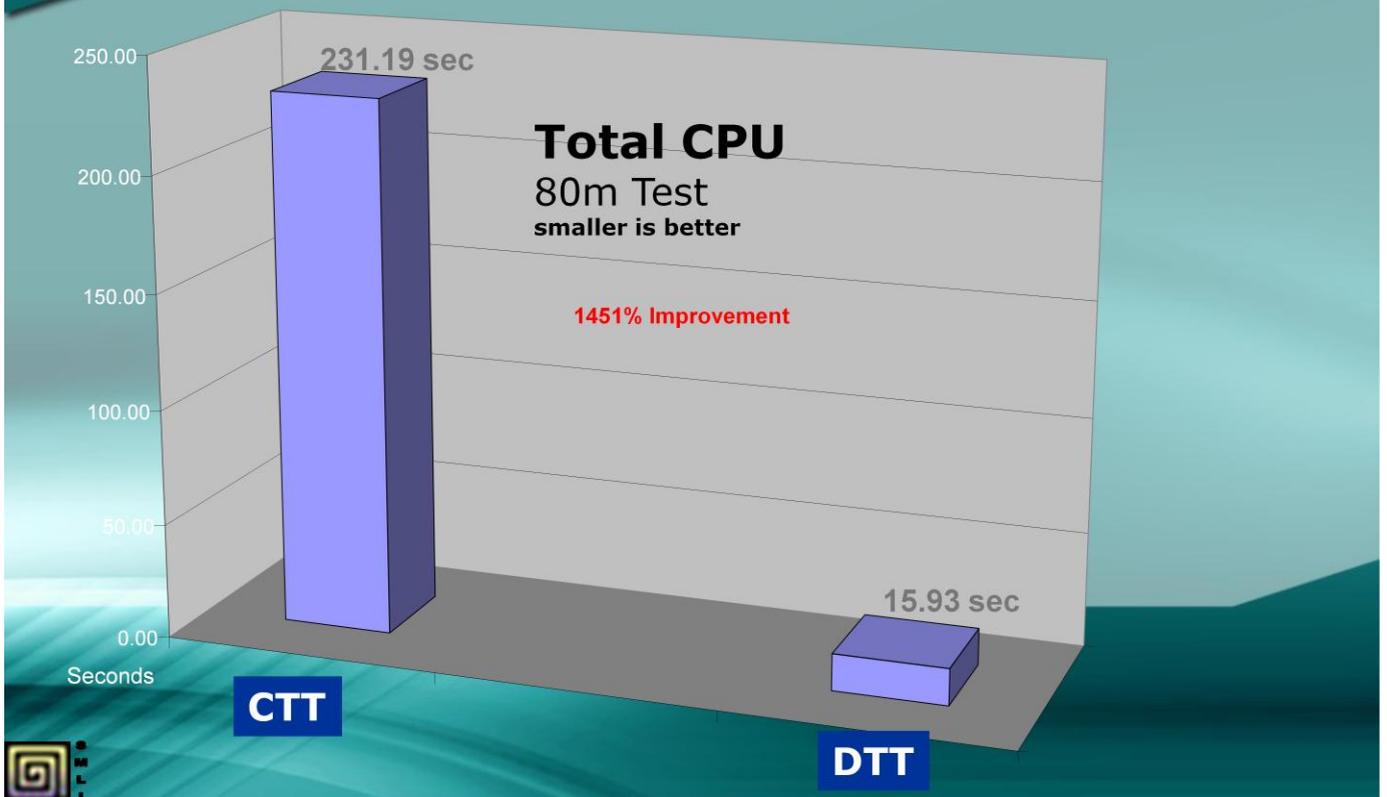
C

## SQL exploitation and proper use



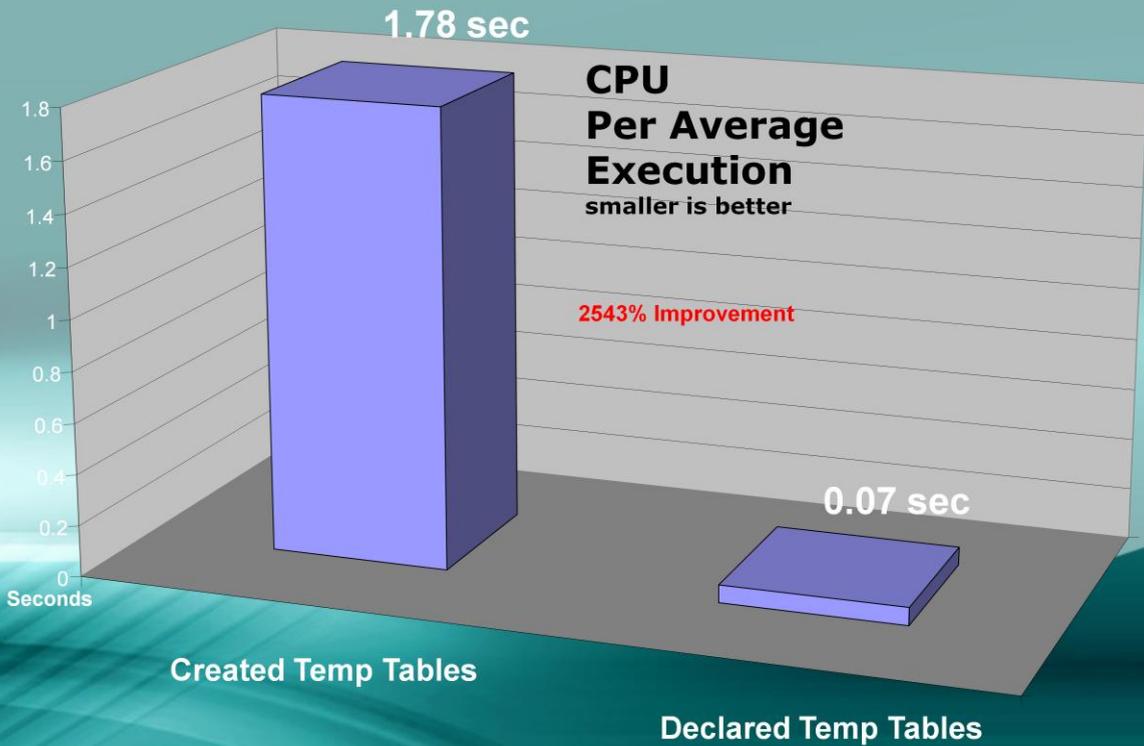
**Notes:** The most resource intensive program was examined in detail and SQL misuse was quickly detected.

# Demand Reduction Initial Refresh



**Notes:** This problem program extracted a tremendous amount of sales activity. The savings depended on the amount of data being extracted. Initial refreshes for 25% of the sales force were reduced 1451%.

# Demand Reduction 6000 Daily Executions



**Notes:** Smaller daily /hourly refreshes were improved 2543%.

# Demand Reduction



## Performance Tuning

**2.8 CPU Hours Per Business Day (Mon.-Fri.)**

**Business Value**



**Notes:** These improvements add up in the course of a day when the program executes 6,000 times a day.

42,750 CPU Minutes Per Year 1 Off of z/OS

712 CPU Hours Per Year 1 Off of z/OS

# Tuning SQL

## Non-Optimal Index Design



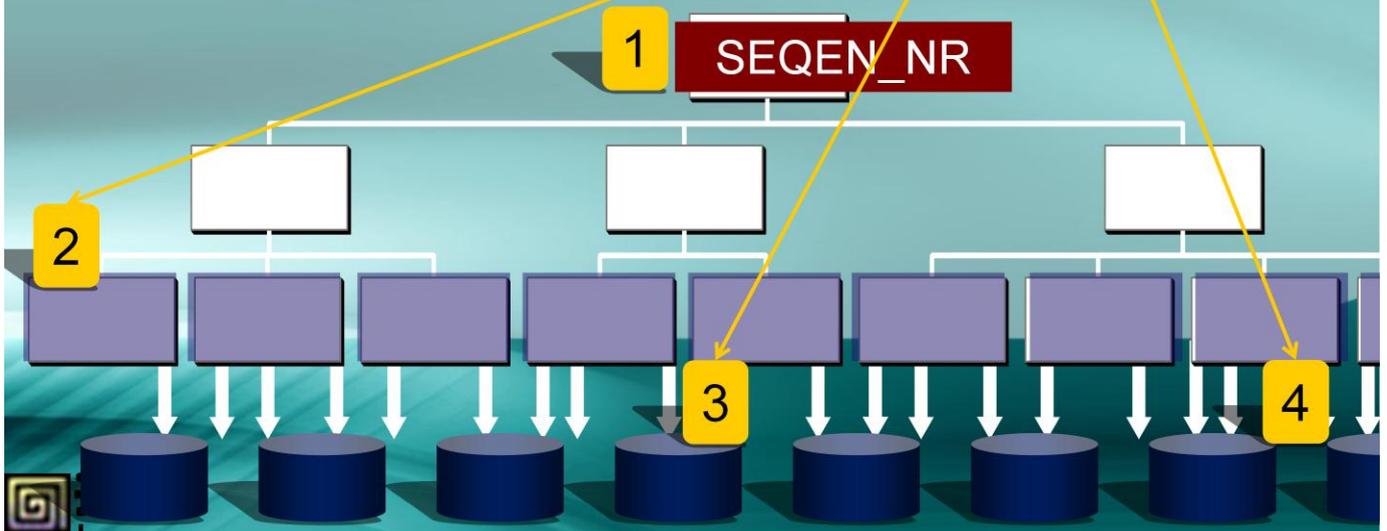
25

**Notes:**

# Four Points of Filtering – DB2

1. Indexable Stage 1 Probe
2. Stage 1 Index Filtering
3. Stage 1 Data Filtering
4. Stage 2

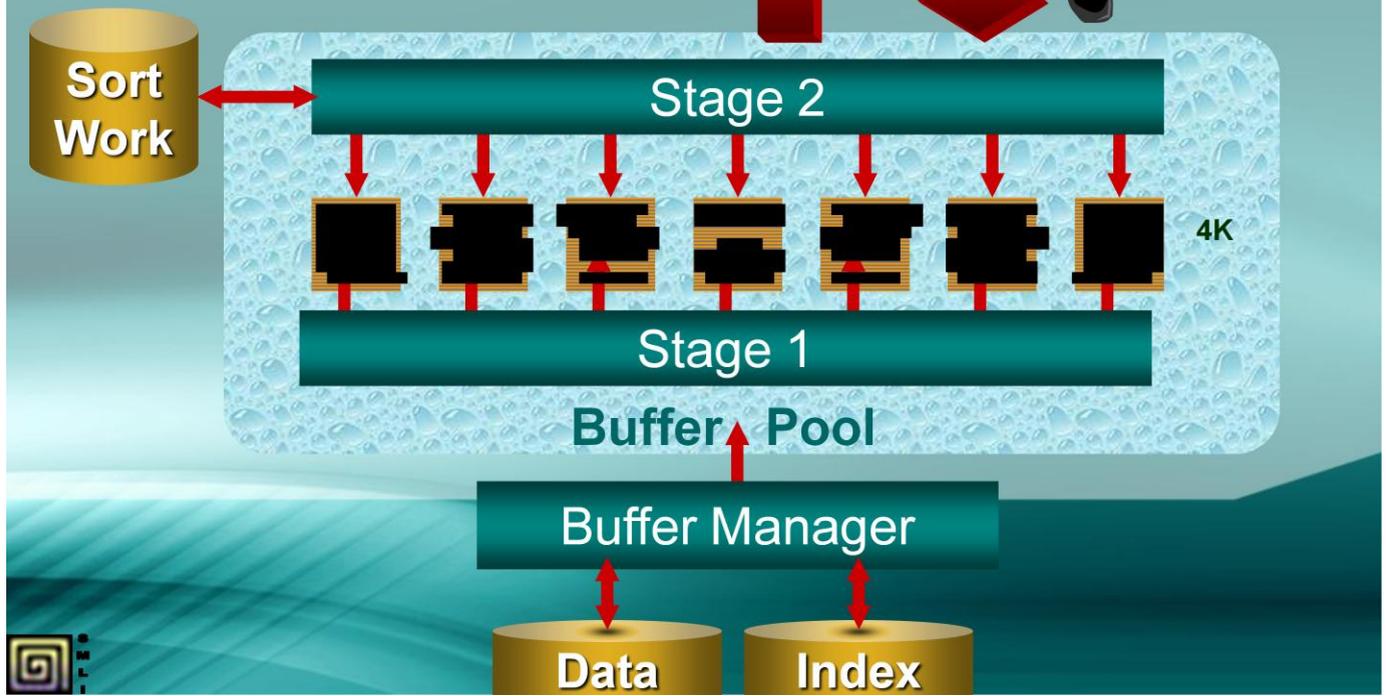
```
WHERE C.LAST_NM LIKE ?  
AND C.SEQEN_NR =  
B.SEQEN_NR  
AND CASE C.COL9 WHEN 'X'  
THEN :hv END ) = 'ABCDE'
```



**Notes:** The first filter to be applied is LAST\_NM LIKE at the 3<sup>rd</sup> point of filtering (after every index page and ever data page was retrieved).

# Search Filters Delayed

Causes heavy buffer pool foot print

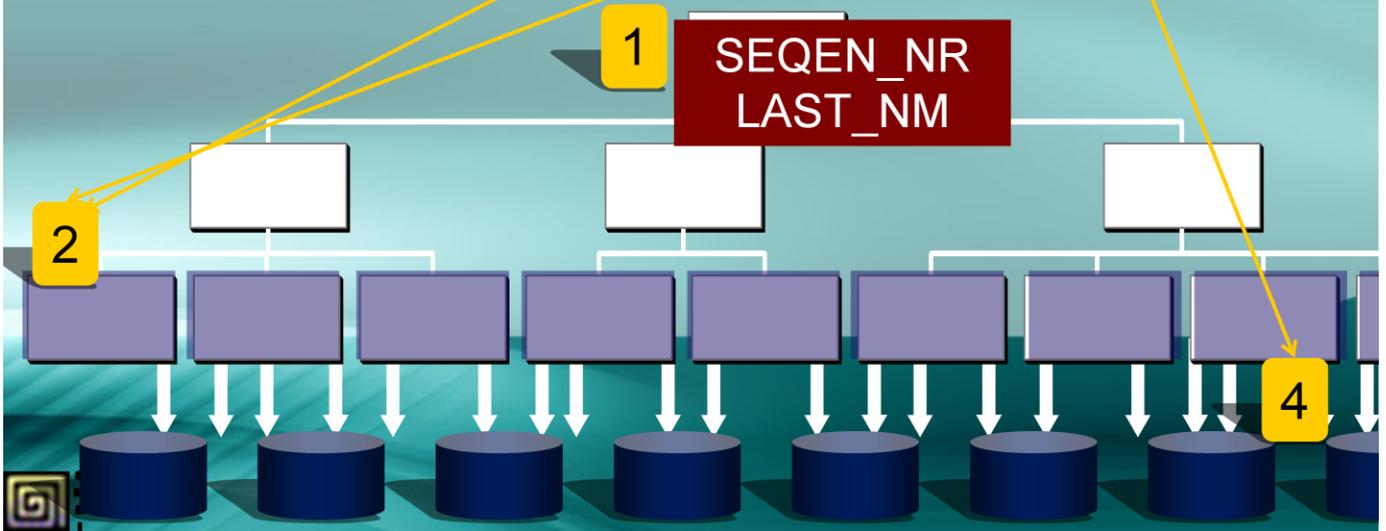


**Notes:** The more pages brought in to the Buffer Pool the larger the foot print. This increases contention and reduces through put.

# Add LAST\_NM to Index

1. Indexable Stage 1 Probe
2. Stage 1 Index Filtering
3. Stage 1 Data Filtering
4. Stage 2

```
WHERE C.LAST_NM LIKE ?  
AND C.SEQEN_NR =  
B.SEQEN_NR  
AND CASE C.COL9 WHEN 'X'  
THEN :hv END ) = 'ABCDE'
```



**Notes:** The first filter to be applied is LAST\_NM LIKE at the 3<sup>rd</sup> point of filtering (after every index page and ever data page was retrieved).

# Delayed Filtering Case Study

- ◆ Web Application Home Page
  - » All tables **Partitioned** and **Clustered** by surrogate key TABLE%\_SEQEN\_NR



29

**Notes:** Architecting web pages to preload all services has a cost but some times can be justified.

# Web Page Details

- 9 Services execute every time
- Dynamic SQL
- Fetches entire result for every service for maximum viewing and sequence agility
- Takes significant wall clock time to load during peak season
- Upcoming slides demo execution of P- Search service



**Notes:** The most expensive service, P-Search was examined.

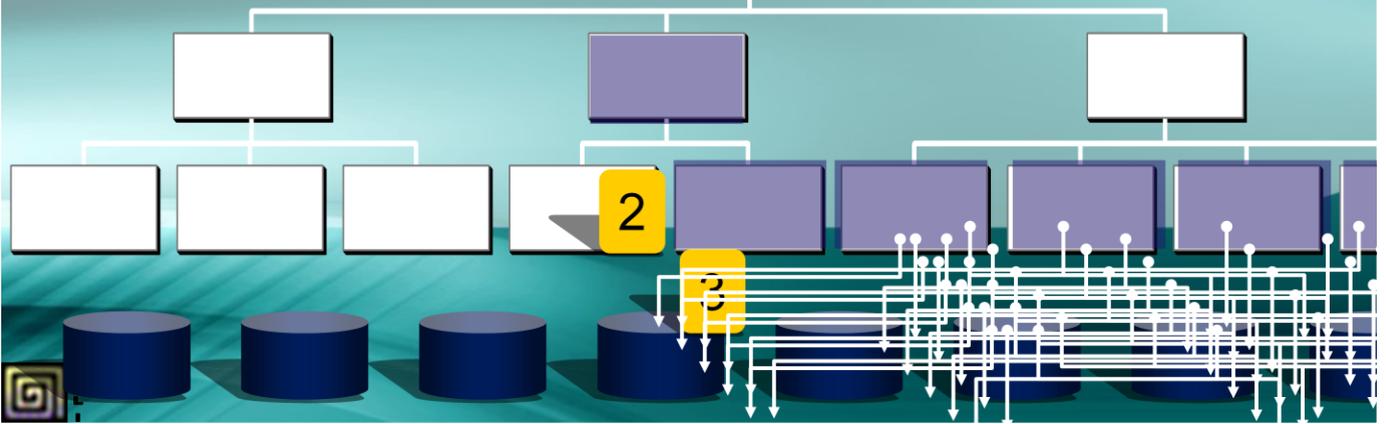
# P--Search starts with TABLEB

**Index – TABLEB4**  
**2 Columns, Nonclustered**

```
FROM TABLE C, TABLE B,  
TABLE A,  
WHERE LAST_NM LIKE ?  
AND C.SEQEN_NR =  
B.SEQEN_NR
```

Type 2 Index

1 LAST\_NM. FIRST\_NM



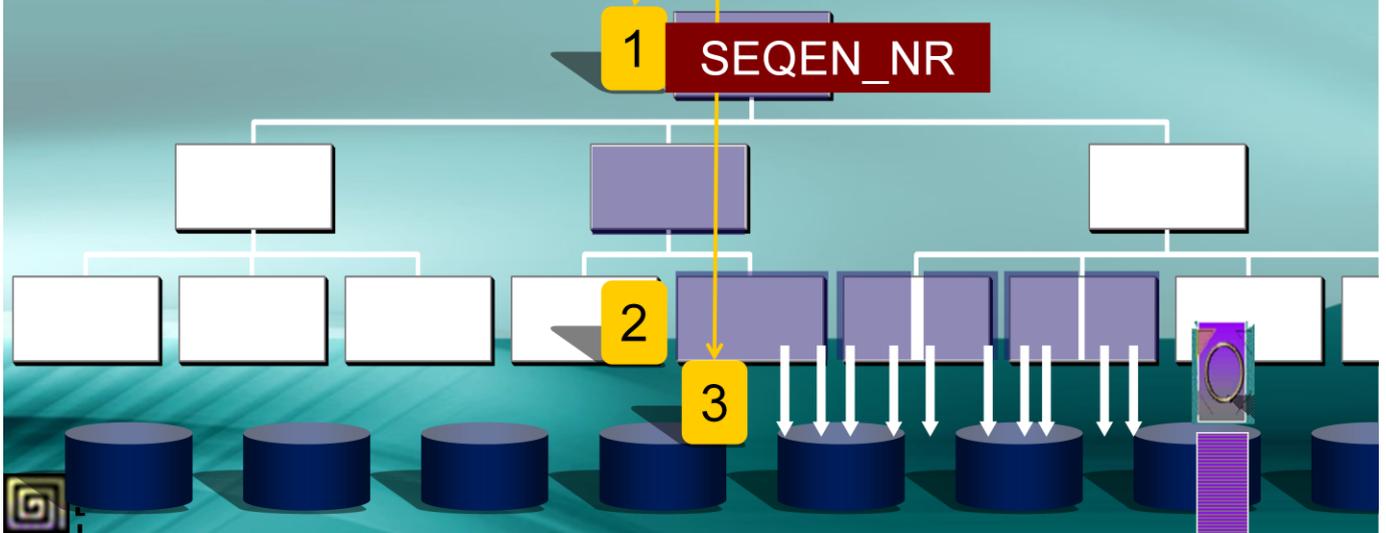
**Notes:** A three table join and DB2 Optimizer starts with TableB and filters LAST\_NM but picks up every SEQEN\_NR using lots of random I/O.



# Joins last to TABLEA

Index  
Clustered –  
TABLEA1  
1 Column

```
FROM TABLE C, TABLE B,  
TABLE A,  
WHERE SEQEN_NR BETWEEN ? AND ?  
AND GUAR_DT >= '2009-04-19'  
AND B.SEQEN_NR = A.SEQEN_NR  
ORDER BY GUAR_DT DESC
```



**Notes:** The last table is joined using Nested Loop because the optimizer chose the join column index instead of the filtering index as no combined index exists and a local filter is supplied using a TABLEC\_SEQEN\_NR BETWEEN ? AND ?. GUAR\_DT is not filtered until the 3<sup>rd</sup> point of filtering and then a sort is issued for the ORDER BY.

# Solutions for P-Search

- ◆ Altered index TABLEI4 added SEQEN\_NR, and SSN\_NR to the end for index only access *(30% CPU reduction per execution)*
- ◆ Recommend altering TABLEI3 to add SEQEN\_NR
- ◆ Recommend adding new index TABLEI5
  - » GUAR\_DT DESC.SEQEN\_NR
- ◆ After solutions are implemented
  - » *9.9 CPU seconds reduced to .2 CPU seconds*



**Notes:** Performance rule violations usually result in increased CPU or I/O, time to fix the mistake, and ultimately, a cost to the business unit.

What will the cost be? Depends on the mistake and the frequency of the mistake.

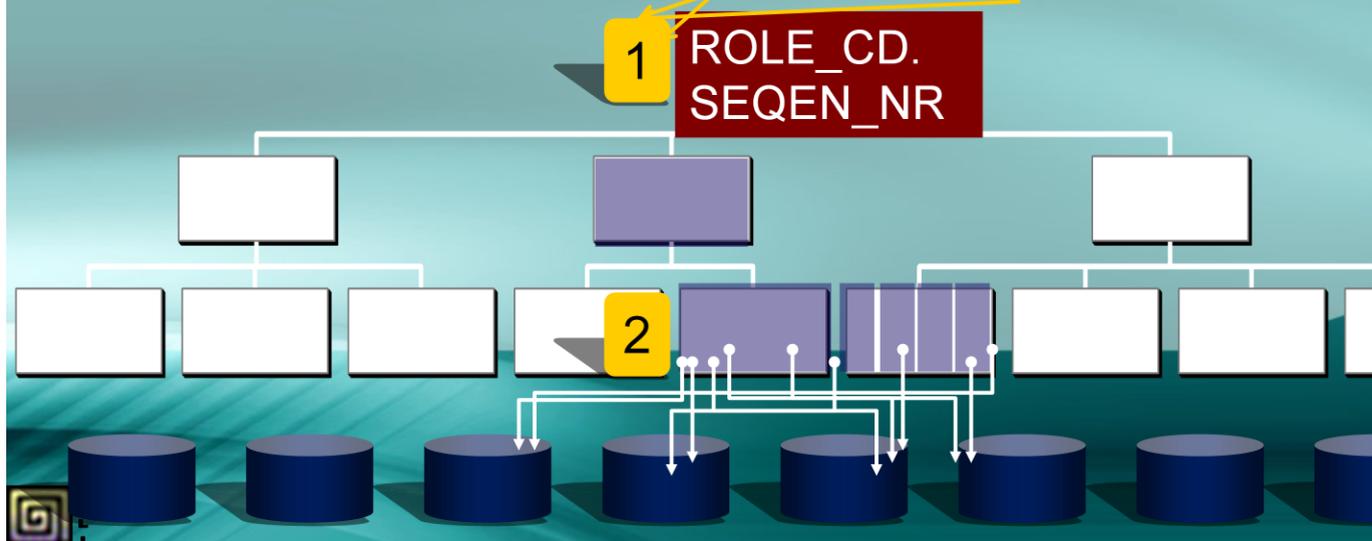
The following presentation will show real case studies of the actual cost of the mistakes.



# Much Less Data Joined

**Index** – TABLEC13  
**2 Columns,**  
**Nonclustered**

FROM TABLE B, TABLE A,  
WHERE **ROLE\_CD = ?** AND  
**TABLEC.SEQEN\_NR BETWEEN ? AND ?**  
AND **C.SEQEN\_NR = B.SEQEN\_NR**

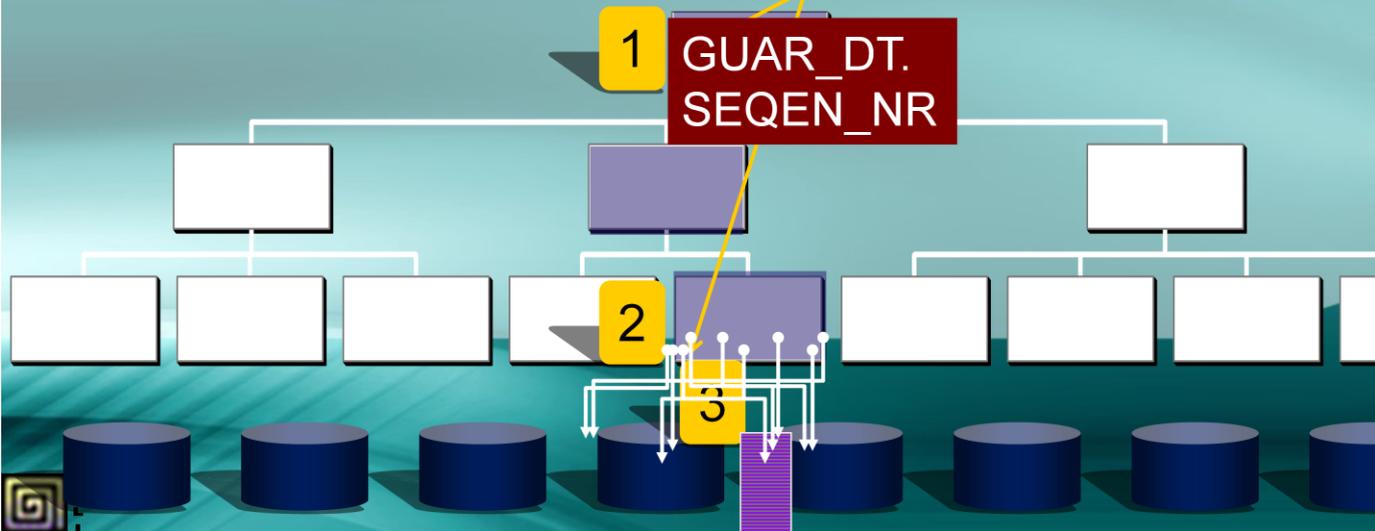


**Notes:** Adding SEQEN\_NR to the filtering index allows the joining to be combined.

# No Sort for ORDER BY

**Index**  
**Nonclustered**  
TABLEA15  
**2 Columns**

```
FROM TABLE C, TABLE B, TABLE A,  
_WHERE TABLEA.SEQEN_NR BETWEEN ? AND ?  
AND GUAR_DT >= '2009-04-19'  
AND B.SEQEN_NR = A.SEQEN_NR  
ORDER BY GUAR_DT DESC
```



**Notes:** By adding SEQEN\_NR to the index the last step combines the joining, filtering and sequencing.

# Cost of SQL Inefficient Use

- ◆ Cost CPU & I/O
- ◆ Cost time to fix
- ◆ Cost \$\$\$\$\$\$\$\$\$

## How much does it cost?

1. Use joins over subqueries when detail row information is required
2. Use subqueries over joins when detail row information is not required
3. Use INNER JOIN over LEFT JOIN when exceptions are not expected or needed
4. Use CREATE GLOBAL TEMPORARY TABLE when 100% data is infrequently accessed
5. Use DECLARE GLOBAL TEMPORARY TABLE with a clustered index when DTT is large and data is frequently accessed



**Notes:** The cost depends on the multiplier. That's when performance and scalability issues become evident.

# Raise Staff Skills to Reduce CPU Consumption

- Use of Powerful New and Old SQL Features
- Don't Misuse SQL
- Don't Use the Wrong Type of Temp Table
- Create Optimal Index Design
- Don't allow Delayed Filtering
- Follow IBM's 23 SQL Performance Rules



**Notes:** Hopefully, you can avoid the mistakes that others have made by following a few performance rules discussed in this presentation.

# IBM Db2 SQL Performance Rules

Table 8-4 Best practices for query design From IBM Data Virtualization Manager for z/OS

Category	Description
Efficient SQL	Do not code mathematics on columns in predicates.
	Sort only on the columns that are needed. No need to ORDER BY EMPNO, LASTNAME when you can ORDER BY EMPNO.
	Watch out for the LIKE predicate. Begins With logic is indexable. Contains is not indexable. Ends With is not indexable.
	Do not code Not Between. Rewrite it as >HV or <HV.
	Use Fetch First XX Rows whenever possible.
	Make sure cardinality statistics exist for all columns in all tables.
	Code Not Exists over Not In. Both are stage 2 predicates but Not Exists typically outperforms the Not In, especially if the list is long.
	When joining two tables the execution is faster if the larger table is on the left side of the join.
	Code WHERE clauses with columns that have unique or good indexes.
	Prioritize WHERE clauses to maximize their effectiveness. First code the WHERE column clauses that reference indexed keys, then the WHERE column clauses that limit the most data, and then the WHERE clauses on all columns that can filter the data further.



## Notes:

# IBM Db2 SQL Performance Rules

Good coding practice	When looking for a small set of records, try to avoid reading the full table by using an index and by providing any possible key values. You can also use more WHERE clauses so that the fetch goes directly to the actual records.
	All Case logic should have an else coded, which eliminates DB2 returning nulls by default if all the Case conditions are not met.
	Stay away from Not logic if possible.
	Minimize the number of times cursors are opened and closed.
	Code stage 1 predicates only. Rewrite any stage 2 predicates.
	Use FOR FETCH ONLY on all read only cursors.
	Reduce the number of rows to process early by using Sub-selects and WHERE predicates.
	Avoid joining two types of columns and lengths when joining two columns of different data types or lengths. One of the columns must be converted to either the type or the length of the other column.
	Limit the use of functions against large amounts of data.



**Notes:**

# IBM Db2 SQL Performance Rules

Category	Description
Reduce impacts to the DVM server	Do not code functions on columns in predicates.
	Minimize the number of times DB2 SQL statements are sent.
	Only select the columns that are needed.
Virtualization	Instead of using multi-level views, try to optimize your SQL queries. Creating views that call other views that call other views can result in joining to the same table multiple times when you only need it once. It creates millions of records in an underlying view where you are interested only in a handful of records.



**Notes:**

# SQL Tuning Confidence Level



0%

100%



Speaker: Sheryl Larsen

Email Address:  
Sheryl.Larsen@kyndryl.com

Phone: (630) 399-3330



**Notes:**