

# Analytics for Cached SQL

Gerald Hodge  
HLS Technologies, Inc.



# Agenda

- Objective
- Data Sources
- Data Extraction
- Analytic Examples
- Tuning Opportunities
- Available SPUIFI code



# Objective

- Collect interval data from the DB2 Cache.
- Use this data to identify:
  - The actual SQL being executed from the Cache after the DB2 Optimization.
  - The actual Explain from the Cache.
  - The SQL performance data from the Cache.

# Data Sources

- We run an extract of data from the Cache at least every hour.
- The extract takes only a few seconds and is non-disruptive.
- The extracted data is processed to provide:
  - The actual SQL executed as a result of the DB2® Optimization.
  - The actual Explain data as present in the Cache for the interval.
  - The measurement data associated with the execution of the SQL
    - The measurement data is cumulative from interval to interval and is modified to show the activity within the interval.



# Analytical Examples

- Identify SQL Optimization differences across time.
- Identify Explain differences
- Identify measurement patterns of interest for both the System and the individual Dynamic SQL
- From the measurement patterns find tuning opportunities
  - One real life example

# Optimization Differences

- It is known that the Dynamic SQL coming from the Optimizer may not match the Dynamic SQL as written
- The Optimizer will try to turn subselects into Joins, etc.
- If there are noticeable differences between interval performance for an SQL Statement then there is a simple SPUFI check to see if the executed SQL statement matches for both intervals

# Explain Differences

- As with the SQL result from the Optimizer the Explain for different intervals can vary
- The Optimizer is influenced by CPU usage, Buffer Pool measurement and other factors at the time of optimization
- This means that the same Dynamic SQL may have a different Access Path in different intervals if the SQL is brought out and back in
- A SPUIFI is provided to check the Explain data between two intervals for a given SQL statement

# Measurement Patterns

- Instead of providing SQL measurements across a complete day information is provided at an hourly interval, this interval can be shortened for finer analysis
- A SPUFI is provided to provide a daily summary for key SQL
- By viewing the interval data it can be determined if specific SQL uses more elapsed time for intervals across the day



# Measurement Patterns: A Real Life Example

- A EU Bank ran one set of SQL in their morning with satisfactory results
- In the late Morning and Afternoon the results were much slower
- There was similar set of SQL coming active during the slow down
- There were Legal reason for the differences between the sets of SQL
- There were Legal / Political reasons that limited the tuning opportunities
- It was identified that the slowdown was a result of contention on a specific index



# Measurement Patterns: A Solution

- A duplicate index was created
- For Auditing / Legal reason the two indexes had to be the same
- The new index was located in a Buffer Pool that was not busy during the business day
- A Hint was made to the first set of SQL to choose the new index
- The result was that the first set of SQL's performance stabilized across the day to what it was in the morning
- The second set of SQL ran as before with a performance improvement



# Available SPUFI Code

- Top nn CPU users
- Top nn Elapsed time users
- CPU / Executions
- Wait Time for Locks
- Wait Time for Sync I/O

