



Tuning the Performance of Dynamic SQL on DB2 z/OS: Using IBM Utilities



Session: IDZ-1031A

Mike Hood, Nationwide
Insurance

IBM Software

Information On Demand **2011**

Agenda

- Dynamic SQL
 - Off Platform Application Calls to DB2 z/OS
 - Driven from Java and other Object Oriented languages
 - Ahdoc Queries
- Poor Performing SQL
 - Multi Table Joins
 - Where END_DT IS NULL
- Tuning Poor Performing SQL
 - Redesign???
 - Reorgs
 - Runstats
 - Optim Data Studio (Explain, Statistics Advisor)



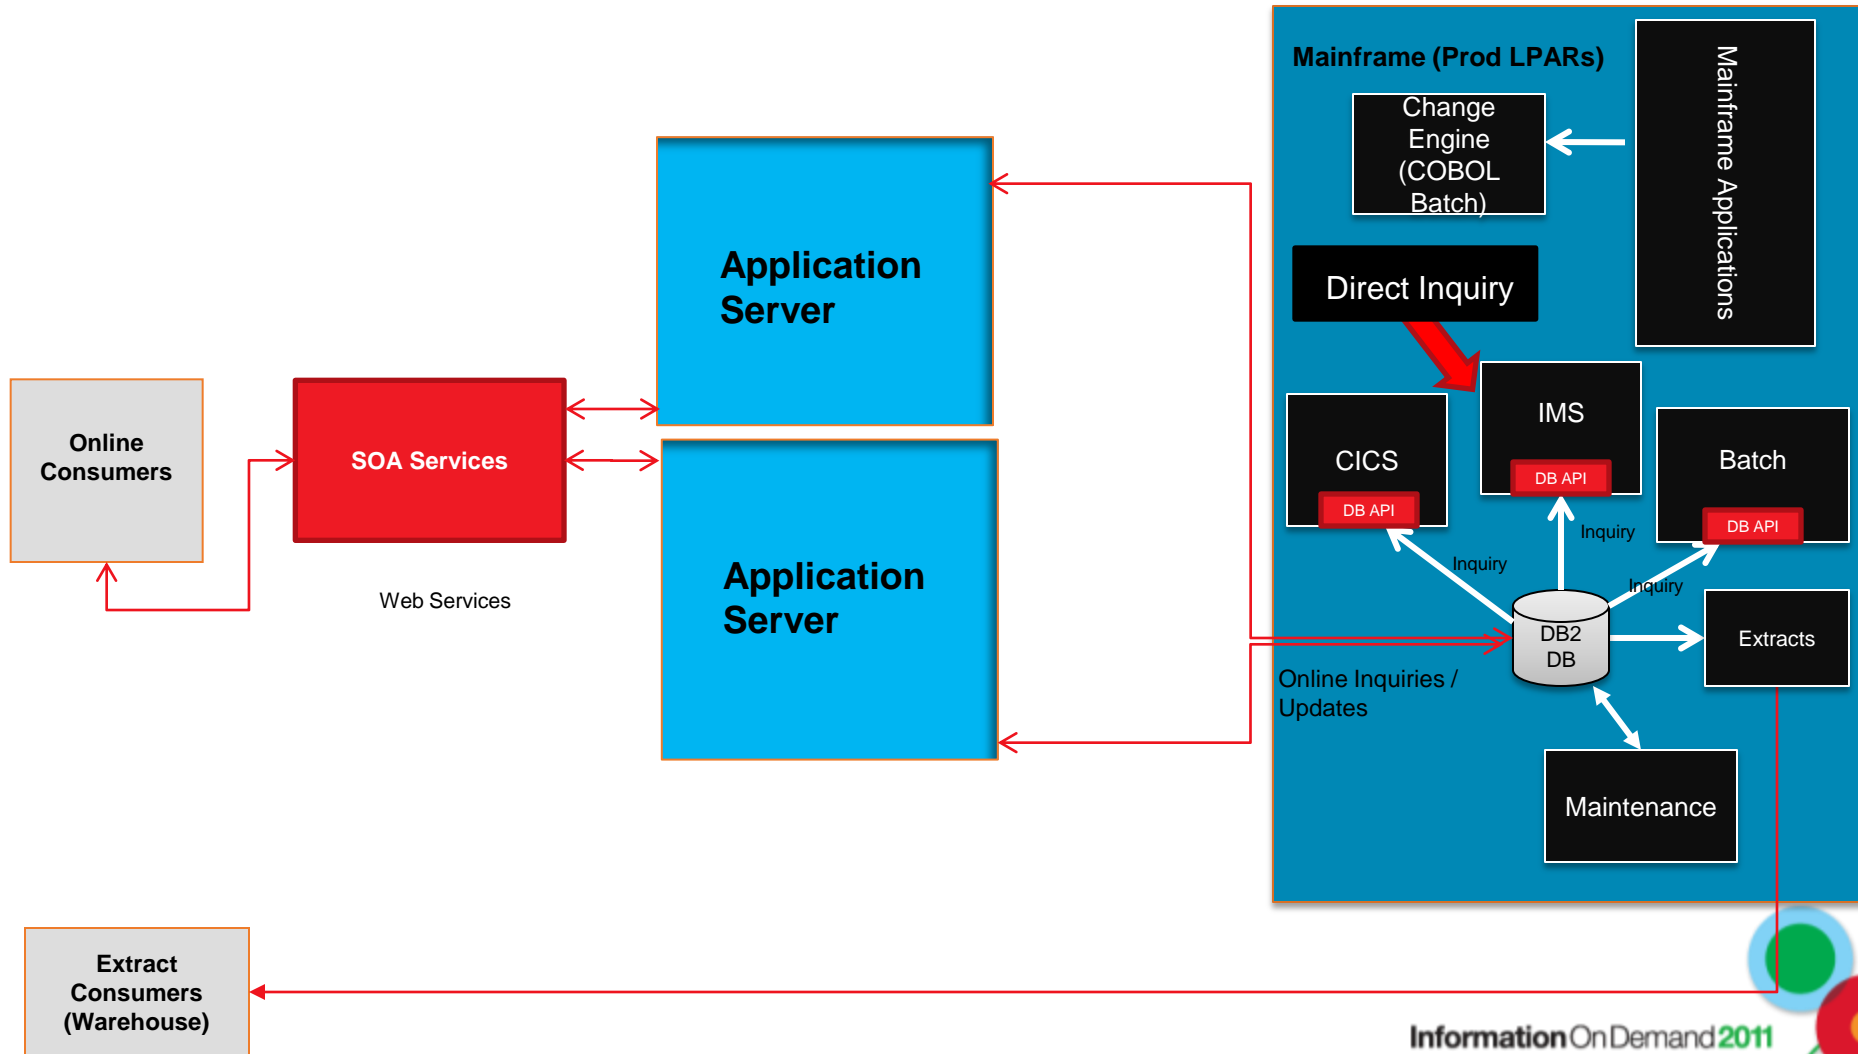




BIO

Mike Hood
Consultant, Database Engineering
Nationwide Insurance
20+ Years in Information Technology
4 Years IMS DBA
7 Years DB2 DBA - Customer File

Typical Dynamic Application Architecture (Conceptual)





Where did all of this crazy code come from?

Large sets of Multi Table Joins
Where WHAT is NULL...
Where IT DOES NOT EXIST???



Dynamic SQL

Where does it come from?

- SQL created by Java using code generators such as Hibernate
- Many vendor packages use Java created SQL
- Extension code by in-house Java developers



Dynamic SQL

How did it get on the mainframe?

- DB2 Connect (Plan DISTSERV)
- Via JDBC TYPE 2 or JDBC TYPE 4 drivers
- Packages
 - SYSLH200 Type 2
 - SYSLN200 Type 4
- NOTE: Use JDBC TYPE 4 for best performance for Java based applications – Significant reduction in overhead



Poor Performing SQL

Generated SQL from Java code

- Java Developers code Object Oriented (i.e. Hibernate, Code Generators)
- Developers do not see the SQL
- No correlation between Java code and SQL
- Generates complex SQL
 - Multi Table joins
 - Multi execution of SQL





Tune Poor Performing SQL

Redesign Tables?

Rewrite SQL?

Not Possible?

What's the Solution?

Tune Poor Performing SQL

Alternative?

Tune what you are dealt...

You can effectively tune Dynamic SQL to perform with some very simple steps

Tuning Poor Performing SQL

Some Examples of poor performers

- WHERE END_DATE is NOT NULL
- WHERE END_DATE is NULL
- WHERE LAST_NAME NOT EXISTS
- SELECT CONTACT_ID
FROM CUSTSEARCH, CUSTOMER, CONTACT
WHERE CUSTSEARCH.postal_code IN(?, ?, ?)
AND CUSTSEARCH.x_street_name IN(?, ?, ?)
AND CUSTOMER.BIRTH_DATE = ?
AND CUSTSEARCH.contact_id = PERSON.contact_id
AND CUSTSEARCH.contact_id = CONTACT.contact_id
AND CONTACT.inactive_date IS NULL



Tuning Poor Performing SQL

Reorgs

- Use DSNACCOX to determine if need exists
 - Use RTS
 - Reorg
 - Runstats

Tuning Poor Performing SQL

Runstats

- Use DSNACCOX to determine if need exists
 - Runstats (common with KEYCARD)
 - Current Runstats play a big role in performance!
 - Explain SQL with Data Studio - Visual Explain

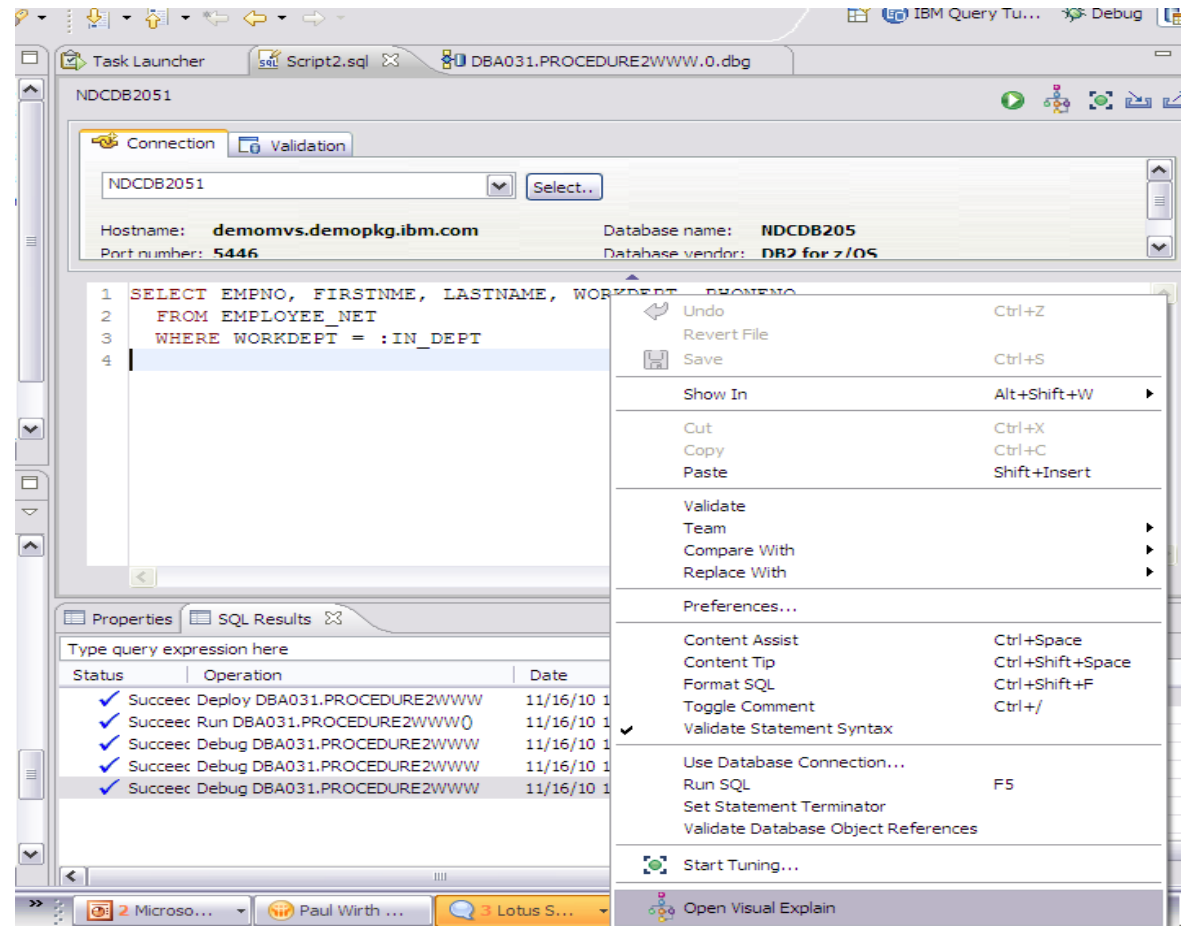
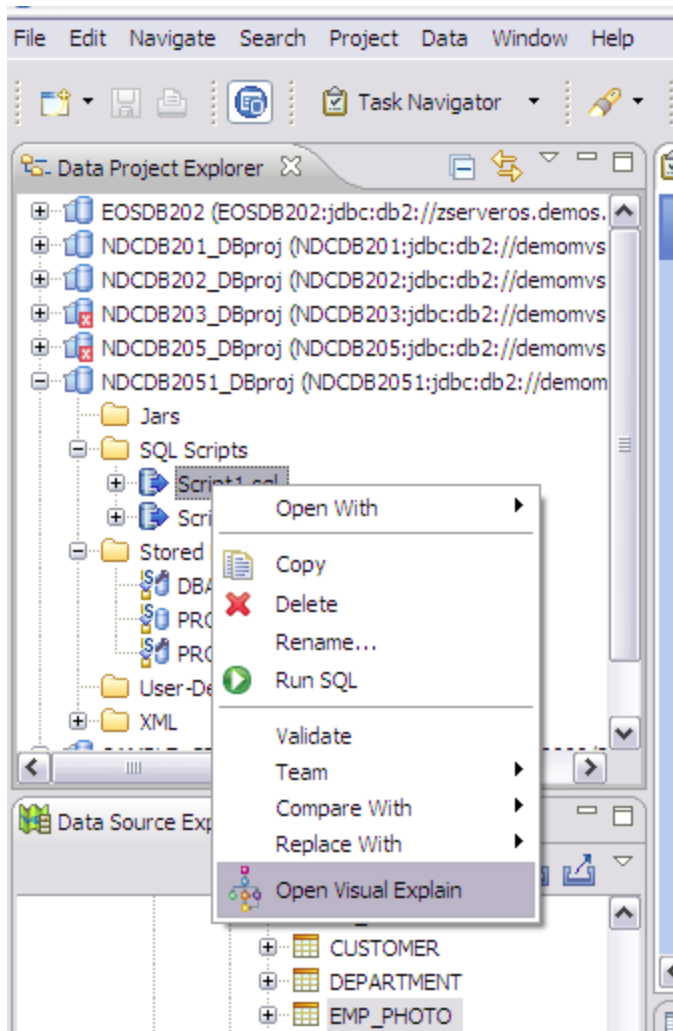


Data Studio

Data Studio is a free tool provided by IBM.

- **Explain facility with Advisors.**
- **Visual Representation of the Access Path**
- **Query Cost Information**
- **Statistics Recommendations**

IBM Data Studio – Visual Explain



IBM Data Studio – Visual Explain

IBM Query Tuning - Query Tuner Workflow Assistant - IBM Data Studio

File Edit Navigate Search Project Run Window Help

Task Launcher NDCDB205_QT_Project/Query Group 1/Query 1 IBM Query Tuning perspective

Query Tuner Workflow Assistant

1. Status

- Text Sources
 - Input Text
 - File
 - SQL Category
 - Exported Workload

2. Capture

- DB2 for z/OS Sources
 - Statement Cache**
 - Catalog Plan or Package
 - QMF
 - QMF HPO
 - SQL Procedure
 - Plan Table
 - Statement Table
 - Function Table
 - View, Trigger, or SQL UDF

3. Manage

4. Invoke

5. Review

6. Compare

DB2 for Linux, UNIX, and Windows Sources

Capture SQL from File

Specify a plain-text file that contains one or more SQL statements and then click Capture to view the captured statements.

Source

File path: <Browse and Select a Text File> Browse Capture

Character set: UTF-8

Statement delimiter: ;

Select Statement Cache For Dynamic SQL



IBM Data Studio – Visual Explain – Select Query

Optim Query Tuner

Server Activation
No product license for IBM Optim Query Tuner is activated on the data server. Advanced tuning functions are disabled. [Learn More](#).

Query Text
Specify the query text for query tuner.

```
SELECT a.grouping_id,  
       a.name,  
       a.grouping_tp_cd,  
       a.entity_name,  
       a.description,  
       a.start_dt,  
       a.end_dt,  
       a.last_update_dt,  
       a.last_update_user,  
       a.last_update_tx_id  
FROM   grouping a,  
       groupingAssoc b  
WHERE  a.grouping_id = b.grouping_id  
AND    b.instance_pk = ?
```

EXPLAIN options: Re-capture EXPLAIN information Use existing server EXPLAIN information

Context

EXPLAIN Options		Application Environment	
Query Number:	111	Schema:	MYDB
SQLID:	DBA1	Current Degree:	<database default>
Optimization hint:	<database default>	Current Refresh Age:	<database default>
Group member:	<database default>	Current Maintained Table Types:	<database default>

Use the following stored procedure to issue the EXPLAIN statement:

Schema: MYDB
Procedure schema: SYSPROC
Procedure name: DSNAEXP

Click Tune Query

Tune Query

IBM Data Studio – Visual Explain – Explain Results With Advisors

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

- Query Format:** Shows the original and transformed SQL. The formatted query is:


```
SELECT A.GROUPING_ID
       , A.NAME
       , A.GROUPING_TP_CD
       , A.ENTITY_NAME
       , A.DESCRPTION
       , A.START_DT
       , A.END_DT
       , A.LAST_UPDATE_DT
       , A.LAST_UPDATE_USER
       , A.LAST_UPDATE_TX_ID
```
- Access Plan Graph:** A visual execution plan diagram. A red arrow points to the root node labeled **@QUERY**. Below it is **@QB1 1**, followed by **@NL JOIN 1.0005**. This join node branches into **@TBSCAN 86.8** and **@FETCH 0.9006**. The **@TBSCAN** node leads to **@GROUPING 5007**. The **@FETCH** node leads to **@IXSCAN 1**, which then leads to **@GROUPINGASSOC 246462**.
- Advisor Recommendation Overview:** A table listing recommendations from the Statistics Advisor.

Advisor	Priority	Description
Recommendations		
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect conflicting st...
Statistics Advisor	MEDIUM	Determine the access path again at runtime. This query contains host variables, paramet...
Statistics Advisor	MAINTENANCE	Gather and recollect all of relevant statistics for this query.

Right Click to Show Description

IBM Data Studio – Visual Explain

**View
Description to
obtain costing
information**

**Take note of the
Cost in MS, SU.
Also Cost
Category.**

Selected Node Descriptor: QUERY

Description of Selected Node
Displays information about the node that is highlighted in the diagram.

..... @ query

Attributes

cost_estimation

Name	Value
Type	SELECT
CPU Cost (ms)	5
CPU Cost (su)	182
Cost Category	B
Reason	HOST VARIABLES
Timestamp	2011-08-26 14:30:15.07

Description of the Selected Attribute



Data Studio – Statistics Advisor



- **Much faster than manually deciding on which stats to collect!**

IBM Data Studio – Visual Explain Statistics Advisor

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

- Query Format:** Shows the original and transformed SQL query. The transformed query is: `SELECT A.GROUPING_ID, A.NAME, A.GROUPING_TP_CD, A.ENTITY_NAME, A.DESCRPTION, A.START_DT, A.END_DT, A.LAST_UPDATE_DT, A.LAST_UPDATE_USER, A.LAST_UPDATE_TX_ID`
- Access Plan Graph:** A visual representation of the query execution plan. The root node is `@QUERY`, which leads to `@QB1`, then `@NL JOIN 1.0005`. This join node branches into `@TBSCAN 88.8` and `@FETCH 0.8006`. `@TBSCAN` further branches into `@GROUPING 5007` and `@IXSCAN 1`. `@FETCH` branches into `@IXSCAN 1` and `@GROUPINGASSOC 246452`.
- Advisor Recommendation Overview:** A table listing recommendations from the Statistics Advisor. A red arrow points to the first recommendation, which is highlighted in blue.

Advisor	Priority	Description
Recommendations		
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect conflicting st...
Statistics Advisor	MEDIUM	Determine the access path again at runtime. This query contains host variables, paramet...
Statistics Advisor	MAINTENANCE	Gather and recollect all of relevant statistics for this query.

Click to Show Statistics Needed



IBM Data Studio – Visual Explain

The screenshot shows the IBM Data Studio interface with the following content:

- Top Panel:** Contains two panes. The left pane, titled "RUNSTATS CONT'D of Statements", shows SQL statements for two tables: "SCIBFM27" and "SCIBFM28". The right pane is titled "PREVIOUS RUNSTATS Statement stored in Database". A red arrow points from the text "Statistics Needed By DB2" to the SQL statements in the left pane.
- Middle Panel:** Titled "Table, index, column, and column group details". It contains a "Statistics Advisor Summary Report" with analysis start and end times. Below this is a "Canvas" section showing details for the table "DBCIBNR.GROUPINGASSOC", including its type, cardinality, collection time, and statistics status (OK). It also lists an index "DBCIBNR.I2_GROUPA" with its own statistics, including a conflicting status.
- Bottom Panel:** Titled "Conflicts detail". It lists conflicts between the table "DBCIBNR.GROUPINGASSOC" and two indexes: "DBCIBNR.I2_GROUPA" and "DBCIBNR.U1_GROUPINGASSOC". Each conflict entry compares the table's cardinality (245452.0) with the index's first key or full key cardinality.

Statistics Needed By DB2

IBM Data Studio – Visual Explain



The screenshot displays the IBM Data Studio interface with the 'Access Plan Graph' for a query. The graph shows the following structure:

- Root node: **@QUERY** (highlighted in green)
- Child node: **@QB1** (0)
- Child node: **@NLJOIN** (1)
- Left child of @NLJOIN: **@FETCH** (1)
- Right child of @NLJOIN: **@FETCH** (1)
- Left child of left @FETCH: **@IXSCAN** (1)
- Right child of left @FETCH: **@GROUPINGASSOC** (12333)
- Left child of @IXSCAN: **@M1_GROUPA** (12333)
- Right child of right @FETCH: **@IXSCAN** (1)
- Left child of right @FETCH: **@GROUPING** (5097)
- Left child of right @IXSCAN: **@U1_GROUPING** (5097)

The text **Success!!** is overlaid in red on the right side of the graph.

IBM Data Studio – Visual Explain

After Correct Statistics are Applied

Selected Node Descriptor: QUERY

Description of Selected Node

Displays information about the node that is highlighted in the diagram.

..... @ query

Attributes

cost_estimation

Name	Value
Type	SELECT
CPU Cost (ms)	1
CPU Cost (su)	2
Cost Category	B
Reason	HOST VARIABLES
Timestamp	2011-08-29 21:44:24.55

Description of the Selected Attribute

Note the Cost reduction.



IBM Data Studio – Visual Explain – Missing Statistics

The screenshot shows the IBM Data Studio interface. At the top, a SQL query is displayed in a window titled 'EMPL_WHICH_DEPTS':

```
SELECT EMPLOYEE.EMPNO, EMPLOYEE.FIRSTNAME, EMPLOYEE.MIDINIT, E  
FROM EMPLOYEE_NET AS EMPLOYEE  
WHERE EMPLOYEE.EMPNO = '000110'
```

Below the query, the 'Access Plan Diagram' is visible. The diagram consists of several nodes: a green circle labeled '(1)QUERY', a purple circle labeled '(2)QB1 400', a yellow hexagon labeled '(3)FETCH 400', a yellow hexagon labeled '(4)IXSCAN 400', and a red hexagon labeled '(5)EMPLOYEE_NET_PK 25(default)'. The IXSCAN node is connected to the EMPLOYEE_NET_PK node, which is in turn connected to the FETCH node, which is connected to the QB1 node, which is connected to the QUERY node. A red box highlights the '(5)EMPLOYEE_NET_PK 25(default)' node, and another red box highlights the '(6)EMPLOYEE_NET 10000(default)' node. An orange arrow points from the 'Diagram Overview' section on the left to the main diagram.

The 'Diagram Overview' section on the left contains the following information:

- Basic Information**
Information that identifies the diagram.
Database Platform: ZOS
Database Version: DSN08015
Explain Timestamp: 2009-10-05T09:20:46.81
[View the SQL Statement](#)
[Save Diagram...](#)
- Diagram Overview**
Display the selected diagram overview.
Query

Tip: Red lines around boxes around tables and indexes indicate missing statistics





Tuning the Performance of Dynamic SQL

- **High Performance can be Achieved**
- **Tuning Dynamic SQL to Perform Using Basic Techniques is Possible**
- **Data Studio Visual Explain is your Friend**



Tuning the Performance of Dynamic SQL

- **Identify Poor Performers**
- **Reorganize the Data (If needed)**
- **Use Data Studio to Explain the Query**
- **Apply Statistics Advisor Recommendations**



Questions??



Disclaimers

© Copyright IBM Corporation 2011. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, IMS, DB2, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.

Disclaimer



Thank You!

Your Feedback is Important to Us

- Access your personal session survey list and complete via SmartSite
 - Your smart phone or web browser at: iodsmartsite.com
 - Any SmartSite kiosk onsite
 - Each completed session survey increases your chance to win an Apple iPod Touch with daily drawing sponsored by Alliance Tech

