An Introduction to SQL Tuning for a DB2 for z/OS Environment

Willie Favero
Data Warehouse on System z Swat Team (DB2 SME)

Session #8762
Tuesday, March 1, 2011: 3:00 PM - 4:00 PM

Copyright © 2011 IBM Corporation
All rights reserved

Notices

This information was developed for products and services offered in the U.S.A.

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Copyright © 2011 IBM Corporation
All rights reserved

Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

  http://www.ibm.com/legal/copytrade.shtml

Copyright © 2011 IBM Corporation
All rights reserved

Acknowledgements

- **Terry Purcell** - DB2 for z/OS Query Optimization
  - "Why did the DB2 for z/OS optimizer choose that access path?"
    - IDUG 2010
  - "What's new from the optimizer in DB2 10 for z/OS?"
    - IDUG 2010
  - "What's new for SQL Optimization in IBM DB2 9 for z/OS"
    - IOD 2009
- **Pat Bossman** - DB2 for z/OS Query Optimization
  - "How to tune a query"
    - 2010
- **James Guo** – DB2 for z/OS Performance
  - DB2 9 for z/OS Selected Query Performance Enhancements
    - IDUG 2007
- **John J. Campbell -** DB2 for z/OS Development
  - "DB2 for z/OS Hints & Tips for Robust Defensive Performance"
    - IOD

Copyright © 2011 IBM Corporation
All rights reserved

Standard Caveats

Copyright © 2011 IBM Corporation
All rights reserved

Performance Goals and Measurements

Performance: the manner in which or the efficiency with which something reacts or fulfills its intended purpose.

Performance Goals: were we want to be, what we want to achieve

Performance Goals
should equal your
Business Goals

Performance Measurements: how we determine if we are making our goals

Your Processes
and Your Tools

Copyright © 2011 IBM Corporation
All rights reserved

What Is Tuning?

- Reducing Wait Times Due to:
  - I/O contention, Locks, Logs, etc...
- Avoiding or reduce I/O
- Controlling Virtual Storage Consumption

Question Number one Should Be:

Where Do I Begin?

Copyright © 2011 IBM Corporation
All rights reserved

BTW, When You're talking I/O…

These two terms are different…

- GETPAGE
  - How DB2 satisfies your request for data from a page
- I/O
  - What DB2 uses if page needed is NOT in the buffer pool

Copyright © 2011 IBM Corporation
All rights reserved

Establish Goals

“You can’t start a journey without a destination”

- What is your SQL doing? Or not doing?
- Is it what you want it to do?
- No, what do you want it to do ?
- How fast is Fast?

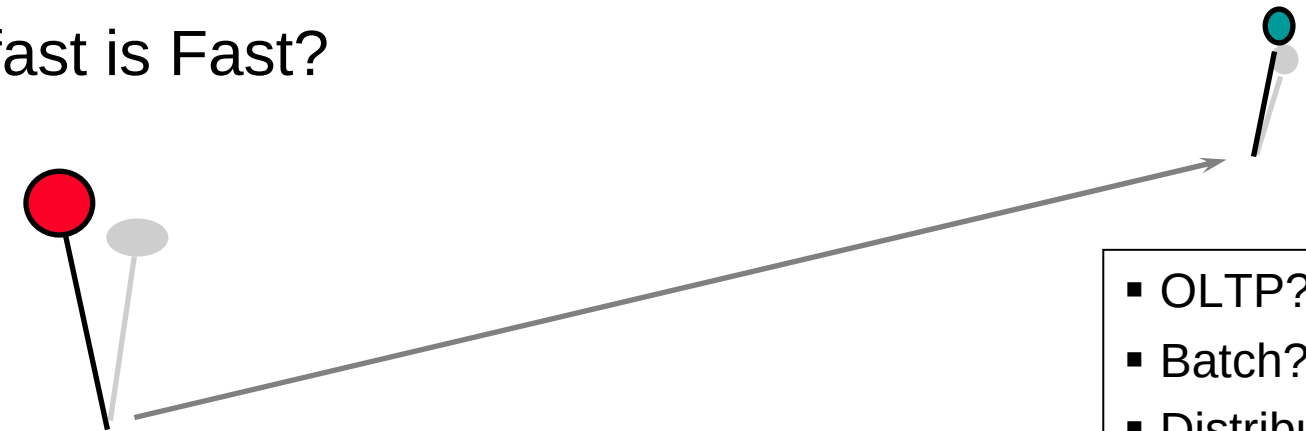

- OLTP?
- Batch?
- Distributed?
- Data Warehouse?

Copyright © 2011 IBM Corporation
All rights reserved

The Process

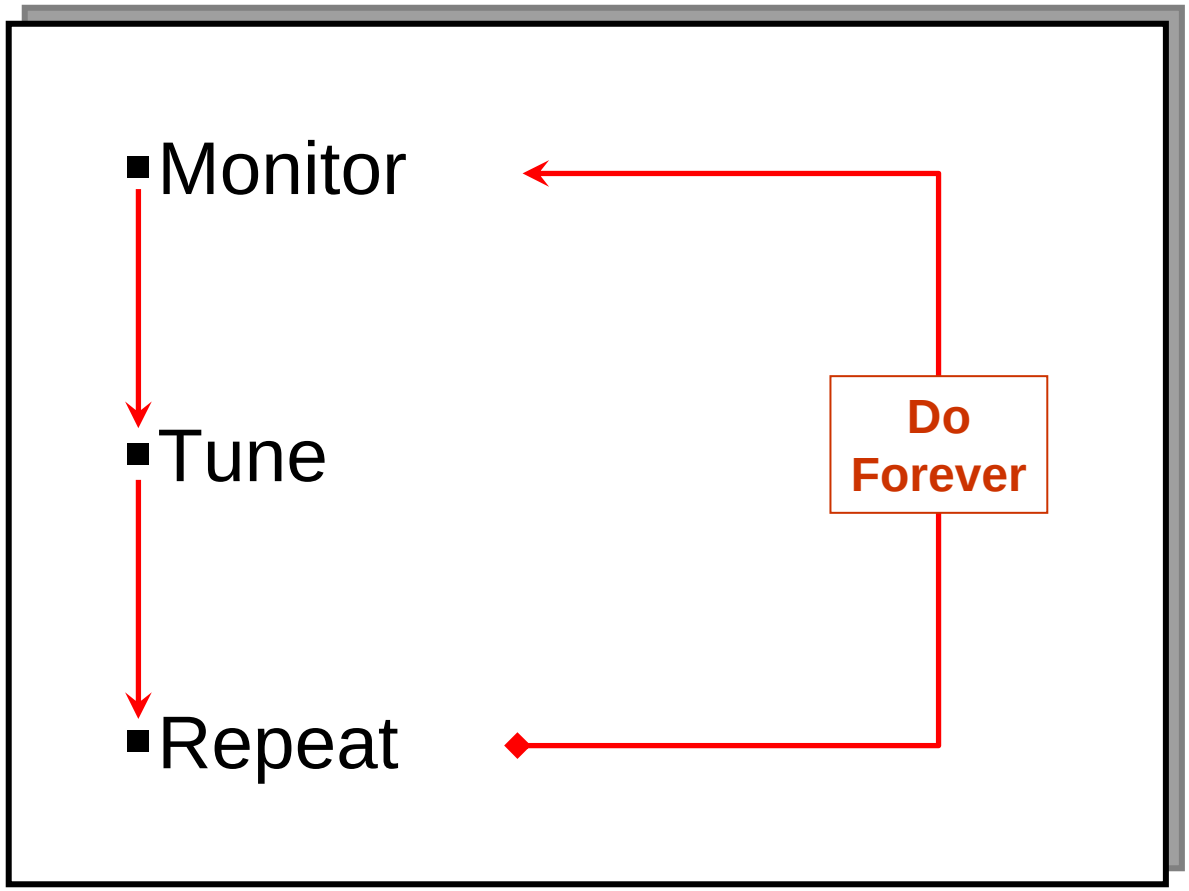

The “Tools” available just keep on getting better

Copyright © 2011 IBM Corporation
All rights reserved

Performance Monitoring Cycle

- Devise a Monitoring Strategy
  - Estimation, Continuous, Periodic, Fire Fighting
- Where to Find Performance Data
- Reviewing Reports
- Determine Constraints/Bottlenecks
- Modify the Application
- Review Reports
  - Were objectives met ?
  - Can the objectives be met ?
  - What are the trade-off ?

Copyright © 2011 IBM Corporation
All rights reserved

What About Traces?

- How are you going to start them?
  - Auto/Manual
- Destination
- Trace Intervals
- Statistics or Accounting Trace Levels
- LUW equivalents

Copyright © 2011 IBM Corporation
All rights reserved

As an Example… (DB2 for z/OS Trace Facilities)

- Accounting
  - Individual Application Performance
  - Class 1, 2, 3, 7, 8
- Statistics
  - DB2 System Wide
  - Class 1, 3, 4, 5, 6
- Audit
  - DB2 Security Control
  - Class 1-8
- All Default to SMF
  - Can go to GTF or IFI
  - SMF Record Types 100, 101, 102
  - Can Be Started by DB2 at startup

Copyright © 2011 IBM Corporation
All rights reserved

Other DB2 Traces

- Performance
  - Monitors Program, Subsystem and Resources
  - Many Classes and IFCIDS
    - Considering tracing specific IFCIDs only
  - Can only be started by command
  - Usual Destination is GTF
- Monitor and Global
  - For Vendors and IBM

Copyright © 2011 IBM Corporation
All rights reserved

Where Could One Spend Their Time Tuning?

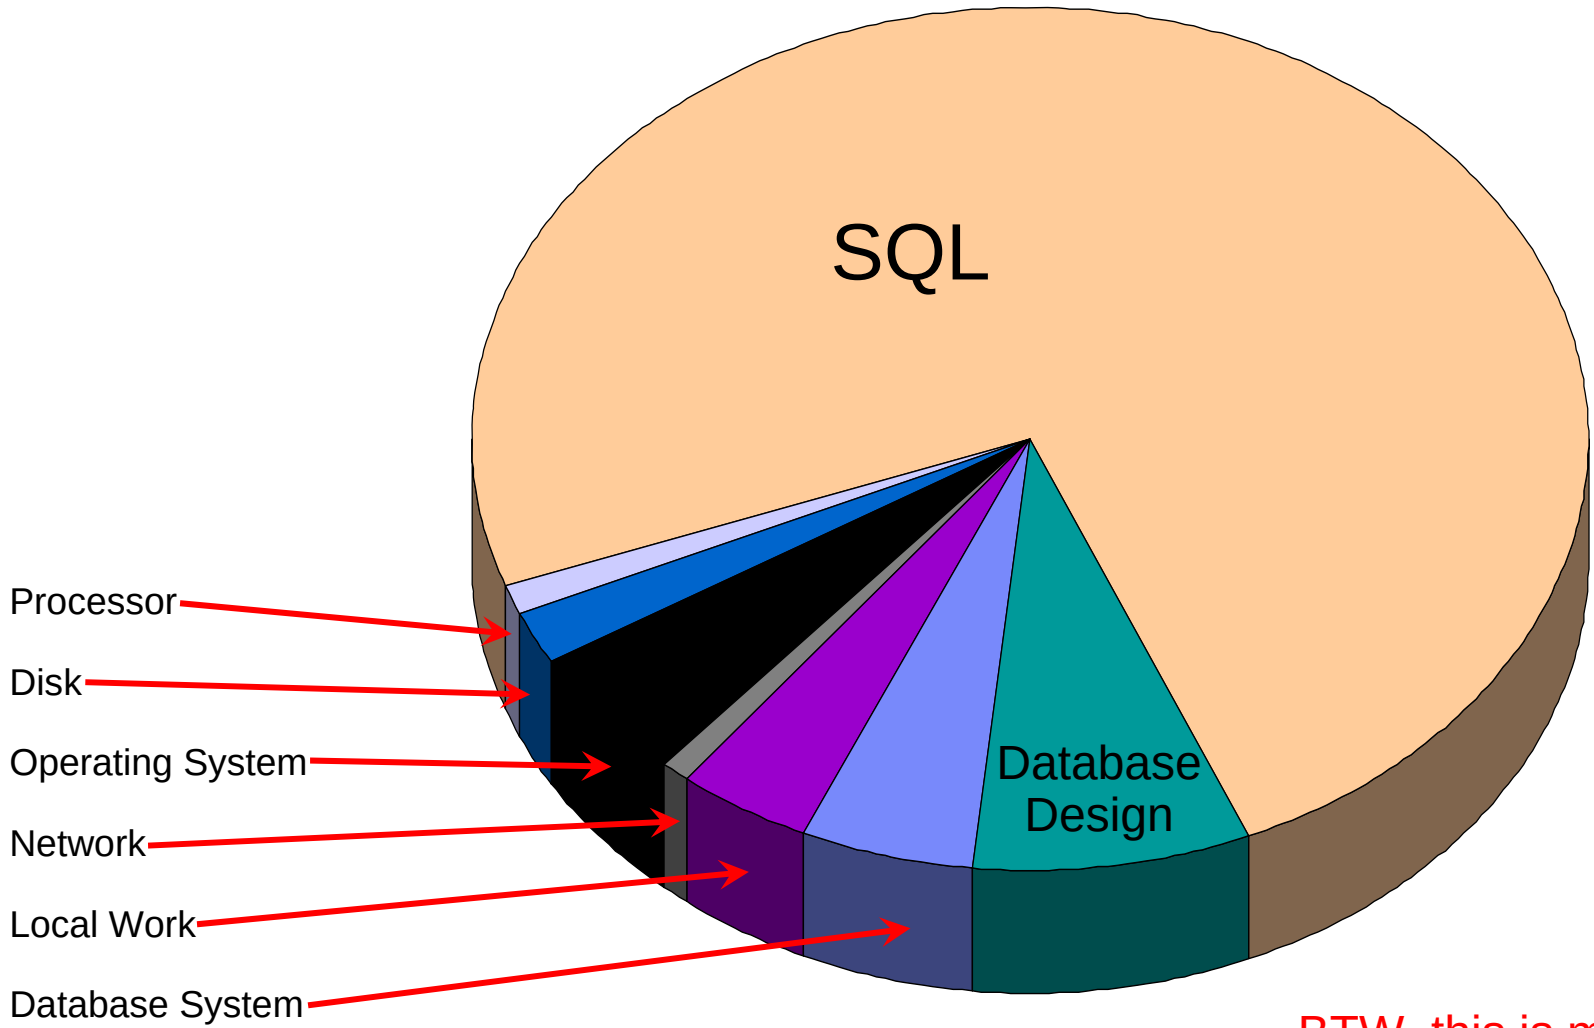

BTW, this is my opinion…

Copyright © 2011 IBM Corporation
All rights reserved

Access Path Selection

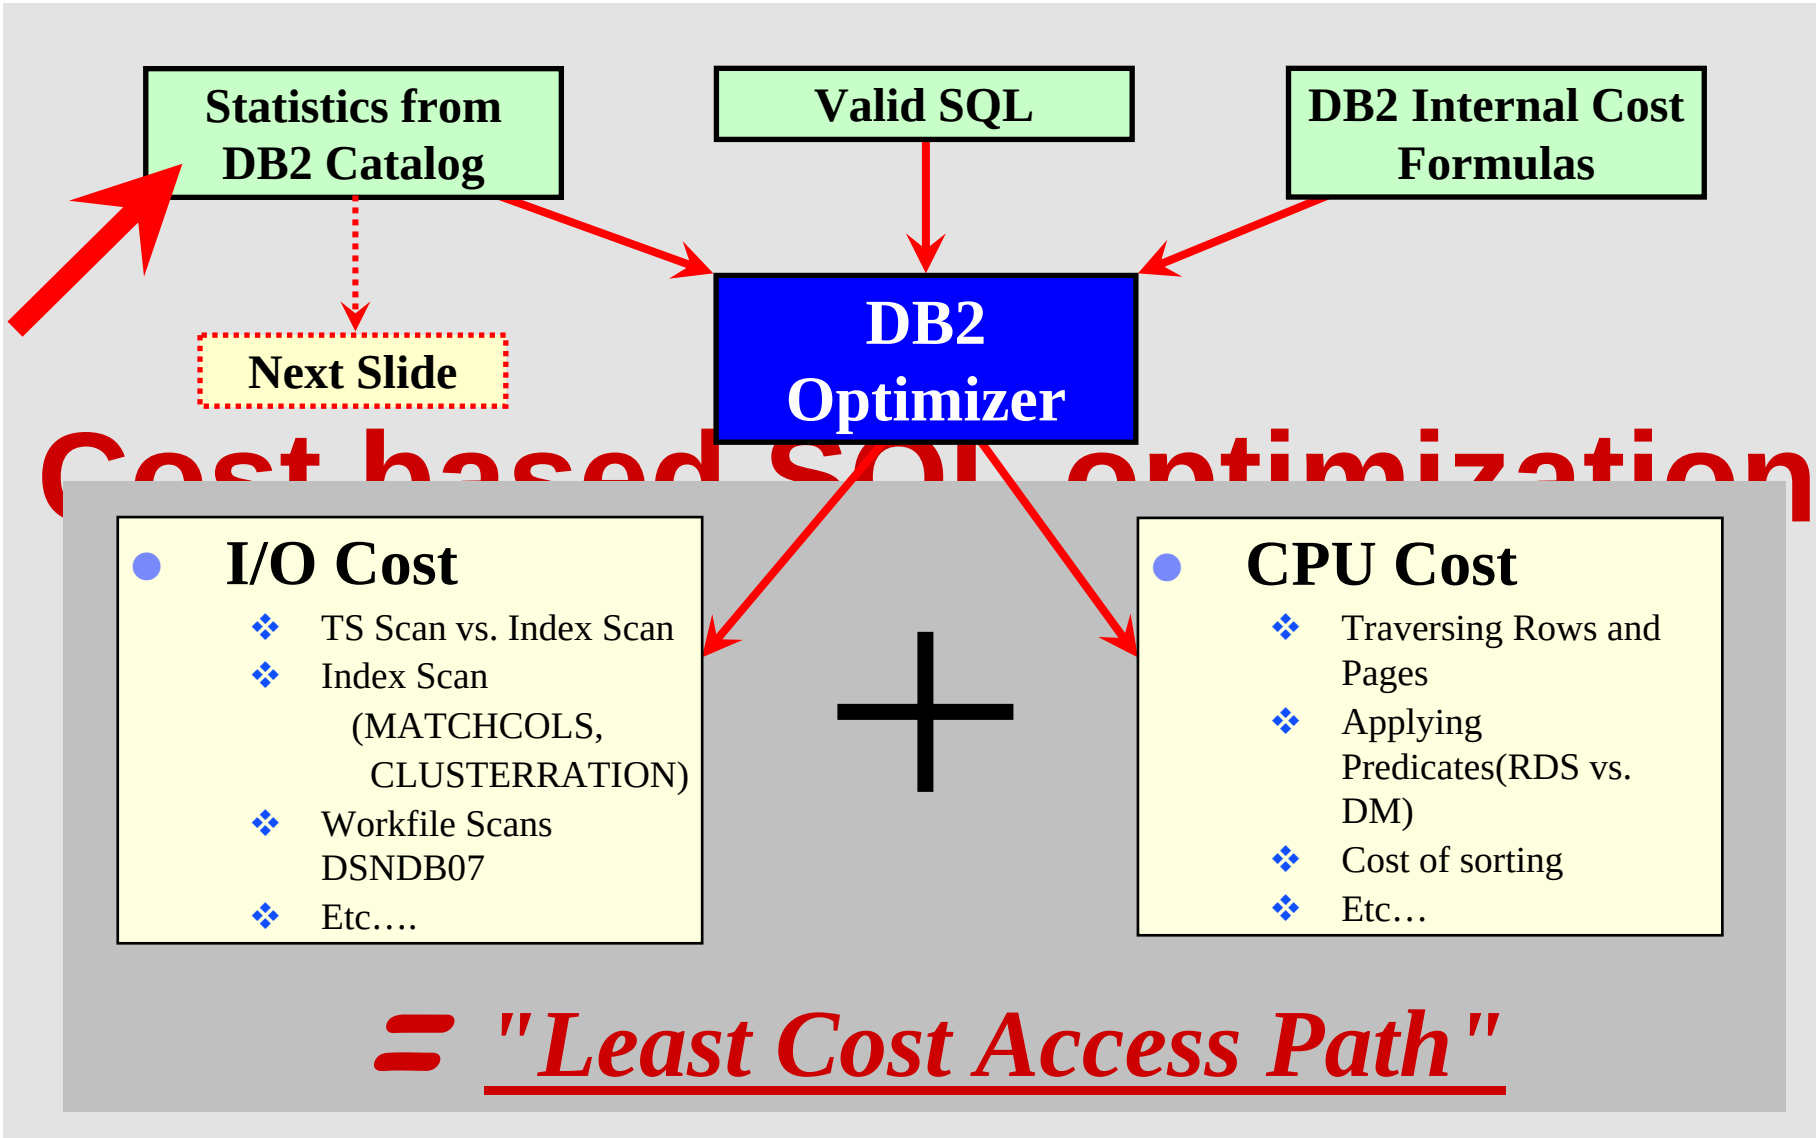

Copyright © 2011 IBM Corporation
All rights reserved

Statistics

- DB2 Catalog tables used for optimization
  - SYSIBM.SYSTABLES
  - SYSIBM.SYSTABLESPACE
  - SYSIBM.SYSTABSTATS
  - SYSIBM.SYSCOLDIST
  - SYSIBM.SYSCOLSTATS
  - SYSIBM.SYSCOLUMNS
  - SYSIBM.SYSINDEXES
  - SYSIBM.SYSINDEXPART
  - SYSIBM.SYSKEYTARGETS
  - SYSIBM.SYSKEYTGTDIST
  - SYSIBM.SYSROUTINES
  - All catalog tables and catalog table columns that affect the optimizer are explained in the "*Performance Monitoring and Tuning Guide*" product publication in Chapter 11

Copyright © 2011 IBM Corporation
All rights reserved

Statistics

- -1 (negative 1) is not a statistic
  - One size does NOT fits all

- RUNSTATS
  - More statistics are not always better
  - Running RUNSTATS does have a cost
  - Consider Statistics Advisor
    - Optimizer Service Center (deprecated in DB2 9)
    - Data Studio (stand alone version only)
    - Optim Query Tuner
    - Optim Query Workload Tuner

Copyright © 2011 IBM Corporation
All rights reserved

Pause for Questions

Copyright © 2011 IBM Corporation
All rights reserved

Unformatted?

```
EXPLAIN PLAN SET QUERYNO = 1 FOR
  SELECT DISTINCT ITEM.ITEM_NBR AS ITEM_NBR,
  ITEM.PRDT_ID, STOREITEM.WK_STRT_DT AS WK_STRT_DT
  ,STOREITEM.DC_ID AS DC_ID FROM PROD.TIPA004_STITM_PROJ
  AS STOREITEM , PROD.TITM001_ITEM AS ITEM WHERE
  ITEM.BUS_UNIT_ID = ‘GS‘ AND ITEM.BUS_UNIT_ID =
  STOREITEM.BUS_UNIT_ID AND ITEM.MJR_CATG_ID = '00754‘
  AND  ITEM.INTMD_CATG_ID = '00043‘ AND ITEM.ITEM_NBR =
  STOREITEM.ITEM_NBR AND ITEM.MJR_CATG_ID =
  STOREITEM.MJR_CATG_ID AND ITEM.INTMD_CATG_ID =
  STOREITEM.INTMD_CATG_ID AND  STOREITEM.RTL_DEPT_NBR =
  1 AND AD_ITEM_FLG = 'Y‘ AND WK_STRT_DT = '2002-02-08';
```


Copyright © 2011 IBM Corporation
All rights reserved

Formatted?

```
EXPLAIN PLAN SET QUERYNO = 1  FOR
 SELECT DISTINCT ITEM.ITEM_NBR AS ITEM_NBR
      , ITEM.PRDT_ID
      , STOREITEM.WK_STRT_DT AS WK_STRT_DT
      , STOREITEM.DC_ID AS DC_ID
 FROM PROD.TIPA004_STITM_PROJ AS STOREITEM
      ,PROD.TITM001_ITEM AS ITEM
 WHERE ITEM.BUS_UNIT_ID = STOREITEM.BUS_UNIT_ID
  AND  ITEM.MJR_CATG_ID = STOREITEM.MJR_CATG_ID
  AND  ITEM.INTMD_CATG_ID = STOREITEM.INTMD_CATG_ID
  AND  ITEM.ITEM_NBR = STOREITEM.ITEM_NBR
  AND  ITEM.BUS_UNIT_ID = ‘GS‘
  AND  ITEM.MJR_CATG_ID = '00754‘
  AND  ITEM.INTMD_CATG_ID = '00043‘
  AND  STOREITEM.AD_ITEM_FLG = 'Y‘
  AND  STOREITEM.RTL_DEPT_NBR = 1
  AND  STOREITEM.WK_STRT_DT = '2002-02-08';
```


Copyright © 2011 IBM Corporation
All rights reserved

Analyzing Query

- Observe “interesting predicates”
  - Optimizer may produce inaccurate filter factor estimate
  - Range predicates with parameter markers
  - Predicates using interesting literals
    - Probable defaults
  - Complex predicates
    - Complex OR expressions
    - Negation predicates
    - Column expressions
    - Non-column expressions

Copyright © 2011 IBM Corporation
All rights reserved

Query Breakdown

```
SELECT  …
FROM    SETL_TRANS S
       ,BRANCH CUST
       ,BRANCH_ADDR A
WHERE S.ADV_ABA_R = ?
AND    S.PROCESS_DT < '9999-12-31‘
AND    S.TYPE_CD  IN ('A', ‘C’, ‘X’)
AND    S.CLR_CYCLE_CD IN ('EOD', 'IMD‘, ‘OPN’)
AND    S.STLMT_DT = ?
AND    S.ACCT_NUM   = CUST.ACCT_NUM
AND    CUST.CUST_EFCT_DT <= ?
AND    CUST.CUST_INACTV_DT >  ?
AND    A.ACCT_NUM = CUST.ACCT_NUM
AND    A.CUST_EFCT_DT    <= ?
AND    A.CUST_INACTV_DT    >  ?
AND    A.ADDR_TYP_CD   = ' '
```


Copyright © 2011 IBM Corporation
All rights reserved

Identify Peculiar Predicates

```
SELECT  …
FROM    SETL_TRANS S
       ,BRANCH CUST
       ,BRANCH_ADDR A
WHERE S.ADV_ABA_R = ?
AND    S.PROCESS_DT < ‘9999-12-31’          ← MAX DATE
AND    S.TYPE_CD  IN ('A', 'C', ‘X’, ‘Z’)
AND    S.CLR_CYCLE_CD IN ('EOD', 'IMD’, ‘OPN’)
AND    S.STLMT_DT = ?
AND    S.ACCT_NUM   = CUST.ACCT_NUM
AND    CUST.CUST_EFCT_DT <= ?               ← Range with marker
AND    CUST.CUST_INACTV_DT >  ?             ← Range with marker
AND    A.ACCT_NUM = CUST.ACCT_NUM
AND    A.CUST_EFCT_DT    <= ?               ← Range with marker
AND    A.CUST_INACTV_DT    >  ?             ← Range with marker
AND    A.ADDR_TYP_CD   = ' ‘                ← COL = blank
```


Copyright © 2011 IBM Corporation
All rights reserved

Why Are They Peculiar?

Predicates with typical default often skewed

```
AND    S.PROCESS_DT < '9999-12-31'      ← MAX DATE
AND    A.ADDR_TYP_CD   = ' '            ← COL = blank
```

Range predicates with parameter markers
- Impossible to estimate without literal

```
AND    CUST.CUST_EFCT_DT <= ?           ← Range with marker
AND    CUST.CUST_INACTV_DT >  ?         ← Range with marker

AND    A.CUST_EFCT_DT    <= ?           ← Range with marker
AND    A.CUST_INACTV_DT    >  ?         ← Range with marker
```


Copyright © 2011 IBM Corporation
All rights reserved

Analyzing Query

- Embed information within statement
  - Table information
    - CARDF
    - NPAGES
  - Column information for predicates
    - Local predicates
    - Join predicates
  - Observe where the filtering is
    - Selectivity of a predicate is relative to table cardinality
- Investigate “suspicious” predicates
  - Determine actual versus estimated filtering
  - If there is a problem, identify options

Copyright © 2011 IBM Corporation
All rights reserved

Embed Statistics

```
SELECT  …
FROM     SETL_TRANS S          CARDF 1,600,254    NPAGES 21,627
        ,BRANCH CUST           CARDF 31,696       NPAGES 1132
        ,BRANCH_ADDR A         CARDF 58,627       NPAGES 2791
WHERE   S.ADV_ABA_R = ?                           COLCARDF 19,712
AND     S.PROCESS_DT < '9999-12-31'               COLCARDF 11
        LOW2KEY 2004-03-24     HIGH2KEY 2004-04-05
AND     S.TYPE_CD  IN ('A', 'C', 'X', 'Z')        COLCARDF 4
AND     S.CLR_CYCLE_CD IN ('EOD', 'IMD', 'OPN')   COLCARDF 3
AND     S.STLMT_DT = ?                            COLCARDF 13
AND     S.ACCT_NUM   = CUST.ACCT_NUM              COLCARDF 15360 / 26,527
AND     CUST.CUST_EFCT_DT <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     CUST.CUST_INACTV_DT >  ?                  COLCARDF 279
        LOW2KEY 2004-03-04     HIGH2KEY 2004-04-07
AND     A.ACCT_NUM = CUST.ACCT_NUM                COLCARDF 26,527 / 26,527
AND     A.CUST_EFCT_DT    <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     A.CUST_INACTV_DT    >  ?                  COLCARDF 274
        LOW2KEY '2004-03-04'   HIGH2KEY '2004-04-07'
AND     A.ADDR_TYP_CD   = ' '                     COLCARDF 5
```


Copyright © 2011 IBM Corporation
All rights reserved

Embed Statistics

```
SELECT  …
FROM     SETL_TRANS S          CARDF 1,600,254    NPAGES 21,627
        ,BRANCH CUST           CARDF 31,696       NPAGES 1132
        ,BRANCH_ADDR A         CARDF 58,627       NPAGES 2791
WHERE   S.ADV_ABA_R = ?                           COLCARDF 19,712
AND     S.PROCESS_DT < '9999-12-31'               COLCARDF 11
        LOW2KEY 2004-03-24     HIGH2KEY 2004-04-05
AND     S.TYPE_CD  IN ('A', 'C', 'X', 'Z')        COLCARDF 4
AND     S.CLR_CYCLE_CD IN ('EOD', 'IMD', 'OPN')   COLCARDF 3
AND     S.STLMT_DT = ?                            COLCARDF 13
AND     S.ACCT_NUM   = CUST.ACCT_NUM              COLCARDF 15360 / 26,527
AND     CUST.CUST_EFCT_DT <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     CUST.CUST_INACTV_DT >  ?                  COLCARDF 279
        LOW2KEY 2004-03-04     HIGH2KEY 2004-04-07
AND     A.ACCT_NUM = CUST.ACCT_NUM                COLCARDF 26,527 / 26,527
AND     A.CUST_EFCT_DT    <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     A.CUST_INACTV_DT    >  ?                  COLCARDF 274
        LOW2KEY '2004-03-04'   HIGH2KEY '2004-04-07'
AND     A.ADDR_TYP_CD   = ' '                     COLCARDF 5
```


Copyright © 2011 IBM Corporation
All rights reserved

Embed Statistics

```
SELECT  …
FROM     SETL_TRANS S          CARDF 1,600,254    NPAGES 21,627
        ,BRANCH CUST           CARDF 31,696       NPAGES 1132
        ,BRANCH_ADDR A         CARDF 58,627       NPAGES 2791
WHERE   S.ADV_ABA_R = ?                           COLCARDF 19,712
AND     S.PROCESS_DT < '9999-12-31'               COLCARDF 11
        LOW2KEY 2004-03-24     HIGH2KEY 2004-04-05
AND     S.TYPE_CD  IN ('A', 'C', 'X', 'Z')        COLCARDF 4
AND     S.CLR_CYCLE_CD IN ('EOD', 'IMD', 'OPN')   COLCARDF 3
AND     S.STLMT_DT = ?                            COLCARDF 13
AND     S.ACCT_NUM   = CUST.ACCT_NUM              COLCARDF 15360 / 26,527
AND     CUST.CUST_EFCT_DT <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     CUST.CUST_INACTV_DT >  ?                  COLCARDF 279
        LOW2KEY 2004-03-04     HIGH2KEY 2004-04-07
AND     A.ACCT_NUM = CUST.ACCT_NUM                COLCARDF 26,527 / 26,527
AND     A.CUST_EFCT_DT    <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     A.CUST_INACTV_DT    >  ?                  COLCARDF 274
        LOW2KEY '2004-03-04'   HIGH2KEY '2004-04-07'
AND     A.ADDR_TYP_CD   = ' '                     COLCARDF 5
```


Copyright © 2011 IBM Corporation
All rights reserved

Embed Statistics

```
SELECT  …
FROM     SETL_TRANS S          CARDF 1,600,254    NPAGES 21,627
        ,BRANCH CUST           CARDF 31,696       NPAGES 1132
        ,BRANCH_ADDR A         CARDF 58,627       NPAGES 2791
WHERE   S.ADV_ABA_R = ?                           COLCARDF 19,712
AND     S.PROCESS_DT < '9999-12-31'               COLCARDF 11
        LOW2KEY 2004-03-24     HIGH2KEY 2004-04-05
AND     S.TYPE_CD  IN ('A', 'C', 'X', 'Z')        COLCARDF 4
AND     S.CLR_CYCLE_CD IN ('EOD', 'IMD', 'OPN')   COLCARDF 3
AND     S.STLMT_DT = ?                            COLCARDF 13
AND     S.ACCT_NUM   = CUST.ACCT_NUM              COLCARDF 15360 / 26,527
AND     CUST.CUST_EFCT_DT <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     CUST.CUST_INACTV_DT >  ?                  COLCARDF 279
        LOW2KEY 2004-03-04     HIGH2KEY 2004-04-07
AND     A.ACCT_NUM = CUST.ACCT_NUM                COLCARDF 26,527 / 26,527
AND     A.CUST_EFCT_DT    <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     A.CUST_INACTV_DT    >  ?                  COLCARDF 274
        LOW2KEY '2004-03-04'   HIGH2KEY '2004-04-07'
AND     A.ADDR_TYP_CD   = ' '                     COLCARDF 5
```


Copyright © 2011 IBM Corporation
All rights reserved

Embed Statistics

```
SELECT  …
FROM     SETL_TRANS S          CARDF 1,600,254    NPAGES 21,627
        ,BRANCH CUST           CARDF 31,696       NPAGES 1132
        ,BRANCH_ADDR A         CARDF 58,627       NPAGES 2791
WHERE   S.ADV_ABA_R = ?                           COLCARDF 19,712
AND     S.PROCESS_DT < '9999-12-31'               COLCARDF 11
        LOW2KEY 2004-03-24     HIGH2KEY 2004-04-05
AND     S.TYPE_CD  IN ('A', 'C', 'X', 'Z')        COLCARDF 4
AND     S.CLR_CYCLE_CD IN ('EOD', 'IMD', 'OPN')   COLCARDF 3
AND     S.STLMT_DT = ?                            COLCARDF 13
AND     S.ACCT_NUM   = CUST.ACCT_NUM              COLCARDF 15360 / 26,527
AND     CUST.CUST_EFCT_DT <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     CUST.CUST_INACTV_DT >  ?                  COLCARDF 279
        LOW2KEY 2004-03-04     HIGH2KEY 2004-04-07
AND     A.ACCT_NUM = CUST.ACCT_NUM                COLCARDF 26,527 / 26,527
AND     A.CUST_EFCT_DT    <= ?                    COLCARDF 2,496
        LOW2KEY 1994-09-02     HIGH2KEY 2004-04-06
AND     A.CUST_INACTV_DT    >  ?                  COLCARDF 274
        LOW2KEY '2004-03-04'   HIGH2KEY '2004-04-07'
AND     A.ADDR_TYP_CD   = ' '                     COLCARDF 5
```


Copyright © 2011 IBM Corporation
All rights reserved

Range Predicate Interpolation

Table 104. Default filter factors for interpolation

COLCARDF	Filter factor for OP	Filter Factor for LIKE / BETWEEN
>= 100,000,000	1 / 10,000	3 / 100,000
>= 10,000,000	1 / 3,000	1 / 10,000
>= 1,000,000	1 / 1,000	3 / 10,000
>= 100,000	1 / 300	1 / 1,000
>= 10,000	1 / 100	3 / 1,000
>= 1,000	1 / 30	1 / 100
>= 100	1 / 10	3 / 100
>= 2	1 / 3	1 / 10
= 1	1	1
>= 0	1 / 3	1 / 10

Note: Op is one of these operators: <, <=, >, >=.

COMMENT: This is DB2's documented guess for an impossible to estimate Filter factor.

Copyright © 2011 IBM Corporation
All rights reserved

Range With Parameter Marker

Error is how different the optimizers DEFAULT estimate is from ACTUAL filtering.

3) AND CUST.CUST_EFCT_DT <= ? **COLCARDF 2,496**
LOW2KEY 1994-09-02 HIGH2KEY 2004-04-06
ESTIMATED FF WITH LITERAL: = 100%
ESTIMATE WITH MARKER: 1/30 = 3% (97% error)

4) AND CUST.CUST_INACTV_DT > ? **COLCARDF 279**
LOW2KEY 2004-03-04 HIGH2KEY 2004-04-07
ESTIMATED FF WITH LITERAL: = 99%
ESTIMATE WITH MARKER: 1/10 = 10% (89% error)

5) AND A.CUST_EFCT_DT <= ? **COLCARDF 2,496**
LOW2KEY 1994-09-02 HIGH2KEY 2004-04-06
ESTIMATED FF WITH LITERAL: = 100%
ESTIMATE WITH MARKER: 1/30 = 3% (97% error)

6) AND A.CUST_INACTV_DT > ? **COLCARDF 274**
LOW2KEY '2004-03-04' HIGH2KEY '2004-04-07'
ESTIMATED FF WITH LITERAL: = 99%
ESTIMATE WITH MARKER: 1/10 = 10% (89% error)

Copyright © 2011 IBM Corporation
All rights reserved

Suspicious Predicate Analysis

- Conclusion
  - The range predicates with parameter markers introduce significant filter factor error. So we should recognize that this filter factor error can cause significant cost estimation problems for the optimizer – possibly resulting in poor access path choice.

Which brings us to parameter markers….

Copyright © 2011 IBM Corporation
All rights reserved

Pause

simply for effect

Copyright © 2011 IBM Corporation
All rights reserved

Dynamic Statement Cache

- Introduced in DB2 V5
- Re-uses SQL and access path
  - If identical SQL string
  - If same user,…
- Avoids full prepare (like a BIND)
- Good programming practice to use parameter marker (?)
  - ? are parameter markers
  - Ensures SQL is always the same
- Not all programs use ?
  - Ruby On Rails generates literals not ?
  - So SQL can not be re-used in Cache

Copyright © 2011 IBM Corporation
All rights reserved

Activate Dynamic Statement Cache

- Dynamic Statement cache must be active
  - Install panel
    - CACHE DYNAMIC SQL field
    - Acceptable values: 5000 to 1048576 KB
    - Initial calculation:

```
EDMSTMTC = (EDMPOOL * 3) + (EDMPOOL/8)
```

    or
  - DSNZPARM
    - DSN6SPRM macro and CACHEDYN keyword

Copyright © 2011 IBM Corporation
All rights reserved

Literal Replacement

- Dynamic SQL with literals can now be re-used in the cache
  - Literals replaced with **&** (similar to parameter markers but not the same)
- To enable either you:-
  - Put CONCENTRATE STATEMENTS WITH LITERALS in the "*attribute-string*" in the PREPARE's **ATTRIBUTES** keyword
  - Or set **LITERALREPLACEMENT** in the ODBC initialization file
  - Or set the keyword **enableLiteralReplacement='YES'** in the JCC Driver
- Lookup Sequence
  - Original SQL with literals is looked up in the cache
  - If not found, literals are replaced and new SQL is looked up in the cache
    - Additional match on literal usability
    - Can only match with SQL stored with same attribute, not parameter marker
  - If not found, new SQL is prepared and stored in the cache

Copyright © 2011 IBM Corporation
All rights reserved

Literal Replacement …

- Example:

```
WHERE ACCOUNT_NUMBER = 123456
```

  - This would be replaced by

```
WHERE ACCOUNT_NUMBER = &
```

- Performance Expectation
  - Using parameter marker still provides best performance
  - Biggest performance gain for small SQL with literals that can now have a cache
  - NOTE: Access path is not optimized to provided literals
    - True for parameter markers/host variables today
    - Need to use REOPT for that purpose

Copyright © 2011 IBM Corporation
All rights reserved

Start Trace for Cache (1 of 2)

- Non-Data-Sharing

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318)
 DEST(SMF)
```

- Data-Sharing

```
-START TRACE(MON) CLASS(1) IFCID(316,317,318)
 DEST(SMF) SCOPE(GROUP)
```

- Verify

```
-DIS TRACE(*)
```


Copyright © 2011 IBM Corporation
All rights reserved

Start Trace for Cache (2 of 2)

- IFCID 316
  - First 60 bytes of SQL statement plus identifying information and statistics
    - http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.xe_db2.doc/ko2rrd20228.htm
- IFCID 317
  - Used in addition to IFCID 316 to obtain the full SQL statement text
    - http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.omegamon.xe_db2.doc/ko2rrd20228.htm
- IFCID 318
  - Acts as a switch for IFCID 316 to collect all available information
    - Stop and Start for new interval

Copyright © 2011 IBM Corporation
All rights reserved

Create Tables Used by EXPLAIN

- Create necessary EXPLAIN table
  - Optimization Service Center (OSC)

    or
  - Optim Query Tuner (was Optimization Expert)

    or
  - DSNTIJOS job in SDSNSAMP

    or
  - Described in the product manuals
    - SQL Reference
    - Performance Monitoring and Tuning Guide (DB2 9 and above)

Copyright © 2011 IBM Corporation
All rights reserved

Get Dynamic Statement Cache

- Snapping the statement cache to cache table
  - EXPLAIN STMTCACHE ALL
    - Can use SPUFI, DSNTEP2, or anything that allows this SQL statement to be executed
    - Data sharing member specific
    - If possible, use SYSADM when issuing SQL statement. Only statements using objects your authid is allowed to access are snapped
    - If DB2 10
      - SQLADM
      - System DBADM

Copyright © 2011 IBM Corporation
All rights reserved

Invalidating Statement Cache

- Execute RUNSTATS utility with the options
  - UPDATE NONE REPORT NO
    - This combination only invalidate statements in the dynamic statement cache without any data access or computation cost
    - Invalidates statements in all data sharing members
- DROP or ALTER the characteristic a statement is dependent
- REORG REBALANCE
- Rebuild index on a cached statement's related table
- Revoking the statements owner's privileges
  - Includes roles

Copyright © 2011 IBM Corporation
All rights reserved

Statement Removal

- Least Recently Used (LRU) algorithm used to remove statements from cache

NOT FIFO

Copyright © 2011 IBM Corporation
All rights reserved

Pause

This time, a new thought
and to catch my breath

Copyright © 2011 IBM Corporation
All rights reserved

REOPT

- V8 REOPT options
  - Dynamic SQL
  - REOPT(NONE, ONCE, ALWAYS)
  - Static SQL
  - REOPT(NONE, ALWAYS)
- V9 Addition for Dynamic SQL
  - Bind option REOPT(AUTO)

Copyright © 2011 IBM Corporation
All rights reserved

Dynamic SQL REOPT - AUTO

- For dynamic SQL with parameter markers
- DB2 will automatically re-optimize the SQL when filtering of one or more of the predicates changes dramatically such that table join sequence or index selection may change
- Some statistics cached to improve performance of runtime check
- Newly generated access path will replace the global statement cache copy
- First optimization is the same as REOPT(ONCE)
- Followed by analysis of the values supplied at each execution of the statement

Copyright © 2011 IBM Corporation
All rights reserved

REOPT(ALWAYS) - REOPT(ONCE)

- REOPT(ALWAYS)
  - DB2 always uses literal values that are provided for parameter markers
- REOPT(ONCE)
  - DB2 reoptimizes cached dynamic SQL statements at execution time for the first execution of the statement based on literal values that are provided for parameter markers

Copyright © 2011 IBM Corporation
All rights reserved

IBM Smart Analytics Optimizer

Capitalizing on the best of relational and the best of columnar databases

What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance-like, add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.

How is it different?

- **Performance:** Unprecedented response times enable 'train of thought' analyses frequently blocked by poor query performance.
- **Integration:** Connects to DB2 through deep integration providing transparency to all applications.
- **Self-managed workloads**: queries are executed in the most efficient way
- **Transparency:** applications connected to DB2, are entirely unaware of the accelerator
- **Simplified administration**: appliance-like hands-free operations, eliminating many database tuning tasks

See my previous session
IBM Smart Analytics Optimizer for DB2 for z/OS

Breakthrough Technology Enabling New Opportunities

Copyright © 2011 IBM Corporation
All rights reserved

Summary

- Prepare a query for analysis
  - Format
  - Annotation
- Perform query analysis
  - Qualified row estimates realistic?
    - Accurately estimate filtering at table level?
    - Accurately estimate filtering for predicates?
  - Evaluate available choices
    - Reference table, index, indexed columns report
    - Develop understanding of “plausible” and “desirable” access paths
    - Examine EXPLAIN output
- Take targeted action
  - Collect additional statistics
  - Provide additional choices
  - Consider REOPT, optimization hints, tricks

Copyright © 2011 IBM Corporation
All rights reserved

More information on zEnterprise

- IBM zEnterprise Announcement Landing Page: ibm.com/systems/zenterprise196
- IBM zEnterprise HW Landing Page: ibm.com/systems/zenterprise196
- IBM zEnterprise Events Landing Page: ibm.com/systems/breakthrough
- IBM Software: ibm.com/software/os/systemz/announcements
- IBM System Storage: ibm.com/systems/storage/product/z.html
- IBM Global Financing: ibm.com/financing/us/lifecycle/acquire/zenterprise/
- Global Technology Services:
  - vanity: ibm.com/services/zenterprise

Copyright © 2011 IBM Corporation
All rights reserved

धन्यवाद

Hindi

多謝

Traditional Chinese

ขอบคุณ

Thai

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكرا

Arabic

Merci

French

Obrigado

Brazilian Portuguese

Bedankt

Nederlands

多谢

Simplified Chinese

Danke

German

நன்றி

Tamil

ありがとうございました

Japanese

감사합니다

Copyright © 2011 IBM Corporation
All rights reserved

Willie Favero

Senior Certified Consulting IT Software Specialist

Dynamic Warehousing on System z Swat Team

IBM Silicon Valley Laboratory

IBM Academic Initiative Ambassador for System z

IBM Certified Database Administrator - DB2 Universal Database V8.1 for z/OS

IBM Certified Database Administrator – DB2 9 for z/OS

IBM Certified System Administrator – DB2 9 for z/OS

IBM Certified Database Administrator – DB2 10 for z/OS

IBM zChampion

wfavero@attglobal.net

Copyright © 2011 IBM Corporation
All rights reserved