



Willie Favero

DB2 SME

Data Warehousing for System z Swat Team

IBM Silicon Valley Laboratory



www.DB2NightShow.com

Friday, October 10, 2014

10:00 am Central

#db2night

Please Note:

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Acknowledgements and Disclaimers:

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. 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. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. 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 the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.


© **Copyright IBM Corporation 2012. All rights reserved.**

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

IBM, the IBM logo, and ibm.com 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.

Possible Topics

- DB2 service levels
 - The IBM Knowledge Center
 - zIIPs and prefetch
 - Statistics
 - New Error Message
 - Sort
 - Buffer Pools
 - DB2 Accelerator Modeling
- 

How Long Has Your DB2 Been Around?

- DB2 9

- GA: 16-Mar-2007
- EOM: 10-Dec-2012
- **EOS**: 27-Jun-2014
 - 6 months, 7 days

- DB2 10

- GA: 22-Oct-2010
 - 3 years, 11 months, 12 days
- EOM: 06-Jul-2015
- **EOS**: 30-Sep-2017

- DB2 11

- GA: 25-Oct-2013
 - 11 months, 9 days

Where Do I Find Information?

- IBM Redbooks
- Google
 - Probably my most used source
- Craig's book... DB2 Developer's Guide
 - Kindle edition
 - I always have my Kindle with me
 - Also have all of my Redbooks on my Kindle
 - Redbooks also now available in EPUB format
- User Group / Conference Presentations
 - Careful using these, not always kept up-to-date
- Friends
- Seminars, webinars, user groups, conferences
- DB2-L list
 - Careful here also, everyone's an expert.

The IBM Knowledge Center

- IBM Knowledge Center home page URL
 - <http://www.ibm.com/support/knowledgecenter/>
- Replaces InfoCenter
- True URLs (great for me when I blog)
- Save collections
- Save searches
- Share (e-mail, Twitter, LinkedIn)

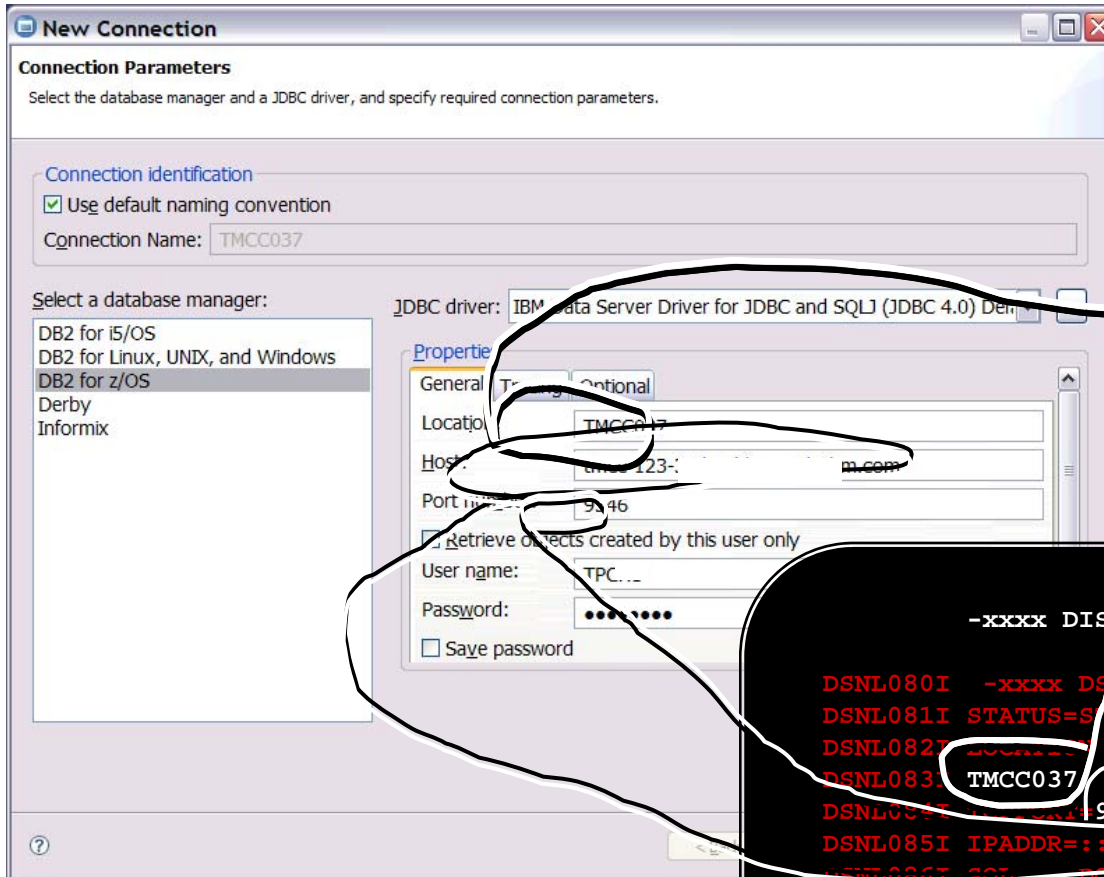
- ***Hello Goodbye* - Knowledge Center now, InfoCenter then...***
 - <http://it.toolbox.com/blogs/db2zos/hello-goodbye-knowledge-center-now-infocenter-then-59144>
- ***Are you an InfoCenter user? Then it's time to change you bookmarks***
 - <http://it.toolbox.com/blogs/db2zos/are-you-an-infocenter-user-then-its-time-to-change-you-bookmarks-59609>

A New Redbook and Redpaper

- ***Reliability and Performance with IBM DB2 Analytics Accelerator V4.1***
 - SG24-8213
 - <http://www.redbooks.ibm.com/abstracts/sg248213.html>
- ***Making Smart Storage Decisions for DB2 in a Flash and SSD World***
 - REDP-5141
 - <http://www.redbooks.ibm.com/redpieces/abstracts/REDP5141.html>
- ***IBM DB2 11 for z/OS Performance Topics*** (no longer draft)
 - SG24-8222
 - <http://www.redbooks.ibm.com/abstracts/sg248222.html>

NEXT

Setting Up a Data Studio Type Client



```
-xxxxx DISPLAY DDF DETAIL

DSNL080I -xxxxx DSNLTDDF DISPLAY DDF REPORT FOLLOWS:
DSNL081I STATUS=STARTD
DSNL082I LUNAME          GENERICCLU
DSNL083I TMCC037        NAME001.DDFTASK  -NONE
DSNL084I 9546          SECPORT=448  REPORT=4461  INAME=-NONE
DSNL085I IPADDR=:tmcc-123-37          .com
DSNL086I 001          DOMAIN=00.00.000
DSNL090I DT=I  CONDBAT= 1000  MEST= 300
DSNL092I ADBAT=   3  QUEDBAT=   0  INADBAT=   0  CONQUED=   0
DSNL093I DSCDBAT=   0  INACONN=   0
DSNL105I CURRENT DDF OPTIONS ARE:
DSNL106I PKGREL = BNDOPT
DSNL099I DSNLTDDF DISPLAY DDF REPORT COMPLETE
***
```

NEXT

zIIPs and Prefetch

- DB2 10
 - buffer pool prefetch, including dynamic prefetch, list prefetch, and sequential prefetch 100% zIIP eligible
- DB2 11
 - GBP (group buffer pool) writes, castout, p-locks and notify/exit processing 100% zIIP eligible
- IIPHONORPRIORITY YES and ZIIPAWMT
- HIPERDISPATCH = YES -- 3.2 msec wait
- HiperDispatch=YES – discretionary goal – no CP
- zAAP on zIIP
 - zEC12 and zBC12
 - zAAP on zIIP, even when no zAAP defined
- Ratio of zIIPs to CPs is 2:1
- IIPHONORPRIORITY
- HIPERDISPATCH = YES -- 3.2 msec wait
- HiperDispatch=YES – discretionary goal – no CP
- ZIIPAWMT

zIIP* Eligible Workload (Sort of Notes)

- DRDA workload increased to 60%
- Utility index maintenance
- **Buffer pool prefetch, which includes dynamic prefetch, list prefetch, and sequential prefetch activities, is 100% zIIP eligible in DB2 10 and DB2 11. (APAR PM30468 reports time under DBM1)**
- Deferred write is zIIP eligible
- Index update parallelism via prefetch
- Portions of RUNSTATS (99% potential)
 - No inline Stats are zIIP eligible
- Native stored procedures initiated from DRDA
- XML parsing

*System z Integrated Information Processor

zIIP Eligible Workload (Sort of Notes)

- DB2 utilities that invoke DFSORT (z/OS 1.10) to sort fixed length records are zIIP eligible if DFSORT selects the memory object sort path. (PK85856, PK85889)
- zAAP on zIIP (z/OS 1.11 + APAR)
- Parallelism
 - Child task
 - Up to 80% eligible (after reaching internal threshold)
 - If coming in from DRDA,
 - Portion of main and child task are eligible
 - Longer running queries, more benefit

II14219: DB2 Z/OS ZIIP EXPLOITATION "SUPPORT USE" INFORMATION

zIIP Eligible Workload


- DB2 utilities
 - Buffer pool prefetch, which includes dynamic prefetch, list prefetch, and sequential prefetch activities, is 100% zIIP eligible in DB2 10 and DB2 11. (APAR PM30468 reports time under DBM1)
- DB2 11 makes all log reads and log writes zIIP eligible

NEXT

Do Not Ignore Statistics

Needed **ONLY** if you plan
to run queries

Populating Statistics



Do I have up-to-date statistics?



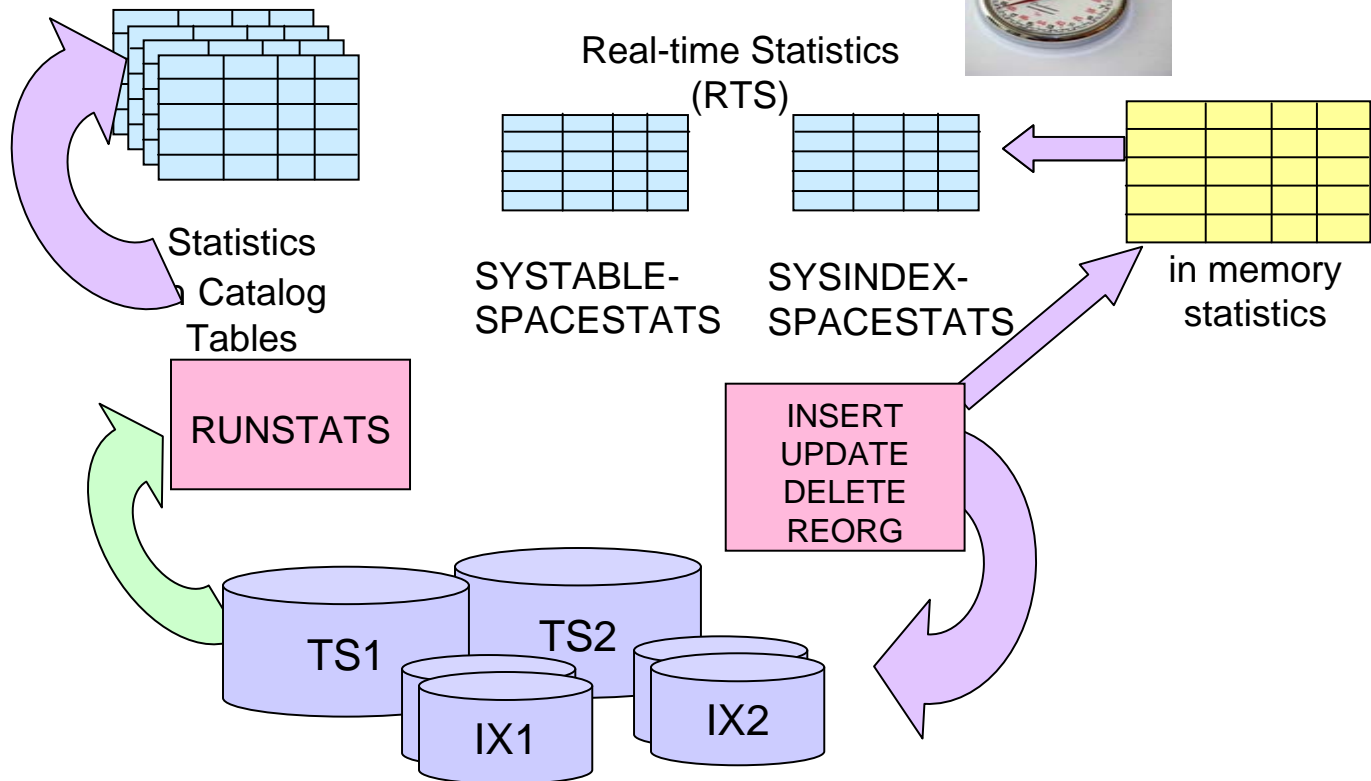
Optimizer

STATSINT
DSNZPARM - minutes



Which statistics to collect?

TABLE(ALL)
INDEX(ALL)
KEYCARD*
HISTOGRAMS
COLGROUP
MCARD
Etc.
Etc.



STATSTIME Column in DB2 Catalog Tables

SYSCOLDIST	SYSCOLDISTSTATS	SYSCOLDIST_HIST	SYSCOLSTATS
SYSCOLUMNS	SYSCOLUMNS_HIST	SYSINDEXES	SYSINDEXES_HIST
SYSINDEXPART	SYSINDEXPART_HIST	SYSINDEXSTATS	SYSINDEXSTATS_HIST
SYSKEYTARGETS	SYSKEYTARGETSTATS	SYSKEYTARGETS_HIST	SYSKEYTGTDIST
SYSKEYTGTDISTSTATS	SYSKEYTGTDIST_HIST	SYSLOBSTATS	SYSLOBSTATS_HIST
SYSSTOGROUP	SYSTABLEPART	SYSTABLEPART_HIST	SYSTABLES
SYSTABLESPACE	SYSTABLES_HIST	SYSTABSTATS	SYSTABSTATS_HIST

Where to start?

- Column – STATSTIME
 - Used in 28 DB2 catalog tables
- If stats have not been collected, column will be 0001-01-01-00.00.00.000000
- Unlike CARDF, COLCARD, NPAGESF, etc..., Never set to -1

Do You Need to Collect Statistics?

- RUNSTATS Utility
 - Use to gather stats for specific query needs
 - Use Data Studio's Statistics Advisor to determine RUNSTATS control cards
 - SAMPLE vs TABLESAMPLE SYSTEM
- Also via the STATISTICS option (inline stats) for
 - LOAD
 - REORG
 - REBUILD
 - Careful - Does not gather all stats you may need
 - i.e. Histogram

More on Statistics Collection

- RUNSTATS needs to be run when the data changes sufficiently to warrant new statistics.
 - For example:
 - REORG of table space or index (use inline stats!)
 - LOAD REPLACE of table space (use inline stats!)
 - After "significant" application changes for the table space or index
 - Periodically (weekly, monthly) except for read only data?
 - Application tracks updates with activity tables?
 - After percentage of pages changed since last RUNSTATS (Check RTS)

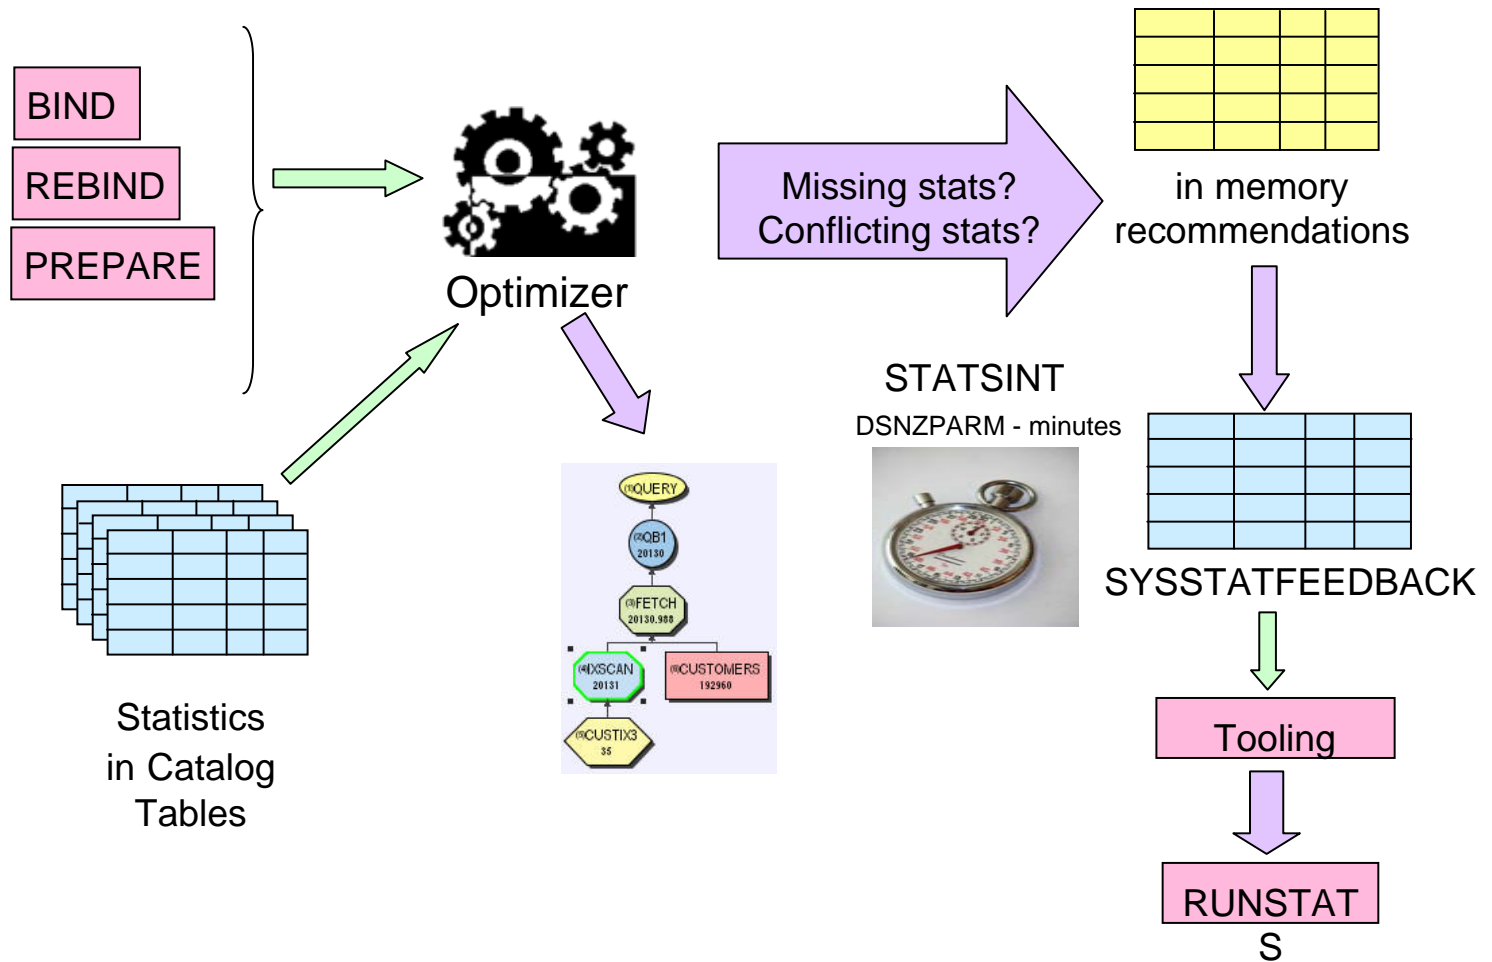
Are Your Stats Correct?

- Are you running adequate statistics control cards?
- Not sure? **Statistics Advisor**
 - aka: Stats Advisor
 - Available in Data Studio (no-charge)
 - Also available in
 - InfoSphere Optim Query Workload Tuner product

Two Possible Approaches

- As you develop your SQL, run Statistics Advisor against the SQL to determine what RUNSTATS control cards are necessary to gather the statistics for that particular SQL statement.
- You can use whatever the normal process is for your statistics gathering, run your SQL statements, and then address specific SQL statements with performance issues by running Stats Advisor against those SQL statements.

DB2 11 – Optimizer externalization of missing statistics



DB2 11 Solution: Optimizer Externalization

- During access path calculation, optimizer will identify missing or conflicting statistics
 - On every BIND, REBIND or PREPARE
 - Recommendations are written asynchronously to SYSIBM.SYSSTATFEEDBACK
 - DB2 also provides statistics recommendations on EXPLAIN
 - Populates DSN_STAT_FEEDBACK synchronously

DB2 11 Solution: Optimizer Externalization

- Contents of SYSSTATFEEDBACK or DSN_STAT_FEEDBACK can be used to generate input to RUNSTATS

Column	Description
TBCREATOR	Creator of the table
TBNAME	Name of the table
IXCREATOR	Creator of the index
IXNAME	Name of the index
COLNAME	Name of the column
NUMCOLUMNS	Number of columns in the column group
COLGROUPCOLNO	Hex representation that identifies the set of columns associated with the statistics. If the statistics are only associated with a single column, the field contains a zero length. Otherwise, the field is an array of SMALLINT column numbers with a dimension equal to the value in NUMCOLUMNS.
TYPE & REASON	Next 2 Slides

- Consistent across all EXPLAIN tables: QUERYNO, APPLNAME, PROGNAME,
- Consistent across almost all EXPLAIN tables: COLLID, GROUP_MEMBER, EXPLAIN_TIME

SYSIBM.SYSSTATFEEDBACK.TYPE

TYPE	Type of Statistic Collected
'C'	Cardinality
'F'	Frequency
'H'	Histogram
'I'	Index
'T'	Table

* This column is defined as CHAR(1)

SYSIBM.SYSSTATFEEDBACK.REASON

REASON	Reason statistics is recommend:
'BASIC'	A basic statistical value for a column table or index is missing.
'KEYCARD'	The cardinalities of index key columns are missing.
'LOWCARD'	The cardinality of the column is a low value, which indicates that data skew is likely.
'NULLABLE'	Distribution statistics are not available for a nullable column.
'DEFAULT'	A predicate references a value that is probably a default value.
'RANGEPRD'	Histogram statistics are not available for a range predicate.
'PARALLEL'	Parallelism could be improved by uniform partitioning of key ranges.
'CONFLICT'	Another statistic conflicts with this statistic.
'COMPPFIX'	Multi-column cardinality statistics are needed for an index compound filter factor.

* This column is defined as CHAR(8):

Externalize In-memory Stats

```
-ACCESS DATABASE(database,...) SPACENAM(space-name,...)  
MODE (STATS)
```

- Externalize in-memory RTS statistics and optimizer recommended statistics
 - In data sharing, externalization is done on all members

Last Word on Stats – Old Statistics

`RUNSTATS TABLESPACE reset-spec`

- Reset all statistics for all tables and related indexes in specified table space
 - RESET ACCESSPATH (optional)
 - No affect on real-time statistics or space stats
 - No stats recovery after RESET completes
 - HISTORY ACCESSPATH
 - Resets SYSTABLES_HIST and SYSINDEXES_HIST
- Not valid for LOBs
- Be sure to run full set of stats afterwards.

NEXT

For Those Who Might be Changing Versions

- Starting at DB2 10 and going forward...
 - DSNY028I
 - csect-name DB2 version ERLY CODE IS NOT COMPATIBLE WITH THIS RELEASE OF DB2
 - 00E80163
 - A DB2 subsystem could not start because the DB2 ERLY code was not at the correct level for the version of that DB2 subsystem.
 - In both cases
 - System action DB2 startup abnormally terminates.

NEXT

DSNZPARAMs – Utility Sort Processing

- IGNSORTN = YES causes DB2 to ignore the SORTNUM keyword if specified. My recommendation is to leave this set to its default NO.
 - If SORTNUM must be specified, setting to YES will prevent SORTNUM's use.
 - Leaving IGNSORTN set to NO gives the option to specifying SORTNUM if necessary.
 - However, it is highly recommended that you NOT use the SORTNUM keyword.
 - Update YES
- Set UTSORTAL to YES, the default, allows DB2 to use real time statistics, always available in DB2 10, to determine the sort work data set sizes.
 - NO forces you to either specify the sort work data sets in the JCL or pre-allocate the data sets.
 - Update YES

More Utility Sort Recommendations

- Dynamic Storage Adjustment (DSA)
 - Initialization value that allows DFSORT to adjust the amount of main storage used by DFSORT (TMAXLIM option).
 - The default for DSA is 64 GB, too low for a DB2 utility with large sorts.
 - 128 GB is the suggested starting value.
 - As high as 512 GB is terabyte sorts tables are expected.

NEXT

Buffer Pools

- Catalog/directory use BP0, BP8K0, BP16K0 and BP32K0
 - Do NOT use for any other purpose
- Minimum of 4 user BPs: user index (4K) and user data (4K) and work files (4K and 32K)
- Don't be afraid to use 8K and 16K buffer pools
 - In many cases can improve rows per page
- Separate dimension tables from fact table
 - If dimension tables are not too large, may be able to pin in pool
 - Same for indexes on dimension tables

Buffer Pools

- Careful with the suggestions to Super Size your pools
 - Always make sure you never exceed real storage
- DB2 10 and above allocates pool space as needed
 - Be careful over allocating pool if you are using PGFIX = YES

Buffer Pools

- Work file (i.e. DSNDB07) buffer pools
 - Need 4k and 32K work files
 - Lots of 32K, maybe more than 4k, monitor
 - VPSEQT = 98 for work file buffer pools
 - If sparse index is used, lower to 90-95
 - Go for LARGE pools if possible
 - 7,000,000 pages large?
- Work files
 - Many and smaller is better than few and large
 - For sort work files, always use zero secondary
 - If using DGTT, make sure you have a few with secondary greater than 0 (zero)

NEXT

Accelerator Modeling

- Provides indicators for possible CPU and elapsed time savings if an IBM DB2 Analytics Accelerator was available
 - It does not require presence of the Accelerator
- DB2 11 or DB2 10
- Controlled by new subsystem parameter **ACCELMODEL** set to **YES** or **NO**
 - If set to YES, DB2 accounting records (IFCIDs 3 and 148) include projected CPU and elapsed time savings
 - CURRENT QUERY ACCELERATION must be set to NONE
 - However, EXPLAIN will still indicate if the query is eligible for acceleration and, if not, the reason why in DSN_STATEMNT_TABLE.REASON
 - Like with any DB2 instrumentation, the new timers need to be formatted and reported by a monitor
- Functionality delivered via two DB2 10 APARs
 - PM90886: aparfix
 - Covers the existing Version 3.1 acceleration capability
 - PM95035: aparfix
 - Adds the new Version 4.1 acceleration capabilities, such as static SQL
 - REBIND needed to enable acceleration modeling

Accelerator Modeling as Reported by OMPE

MEASURED/ELIG TIMES	APPL (CL1)	DB2 (CL2)
-----	-----	-----
ELAPSED TIME	4.830139	4.740227
ELIGIBLE FOR ACCEL	N/A	4.442327
CP CPU TIME	6.337894	6.336111
ELIGIBLE FOR SECP	4.990042	N/A
ELIGIBLE FOR ACCEL	N/A	6.329119
SE CPU TIME	0.000000	0.000000
ELIGIBLE FOR ACCEL	N/A	0.000000

1

2

3

1

Elapsed time that can be significantly reduced because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the reduced elapsed time relates to the parent task only.

2

The part of CPU time spent on general purpose processors that can be saved to a large extent because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the CPU saving includes the parent and all the subordinated parallel tasks.

3

The part of CPU time spent on specialty engine processors that can be saved to a large extent because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the CPU saving includes the parent and all the subordinated parallel tasks.

NEXT

z/OS v2.1 Enhancements used by DB2 11

- 2 GB large pages require z/OS V2R1 and the hardware features of the zEC12
- Enhancing DB2 BACKUP SYSTEM solution (DB2 11 only)
 - Enable recovery of single pageset from DB2 system-level backup even if original volume does not
 - have sufficient space
 - Enable exploitation of FlashCopy consistency group for DB2 BACKUP SYSTEM
 - Enable restore of a page set to a different name
- z/OS DFSMS VSAM RLS for Catalog support (DB2 9 and above)
 - Improved DB2 dataset open/close performance
- z/OS DFSMS Storage Tiers (DB2 10 and above)
 - Optimizes disk placement on SSD and HDD
 - DB2 10 and 11 provide ability to issue STOP DB AT COMMIT synchronously
- Enhanced WLM managed buffer pools (DB2 9 and above)
 - z/OS 2.1 the first version that enables a buffer pool to shrink
- XES GBP write-around support (DB2 11 only)
 - CFLEVEL 17 or above (retrofit to z/OS 1.13 with OA37550)

z/OS v2.1 Enhancements used by DB2 11

- Reduced DRDA message latency and DB2 CPU reduction (DB2 11 only)
 - TCP/IP synchronous receive (retrofit to z/OS 1.13 with OA39810)
- Reduced DRDA message latency and DB2 CPU reduction (DB2 9 and above)
 - CommServer SMC-R support for RDMA over ethernet (RoCE), Network Adapter.
 - Requires zEC12 or zBC12, z/OS to z/OS only for now
- 2G page frame support (DB2 11 only)
 - Requires zEC12 or zBC12
 - Retrofit to z/OS 1.13 – OA40967 and z/OS V1R13 RSM Enablement Offering
- Pageable 1M frame support (DB2 10 and 11)
 - Requires zEC12 or zBC12 (retrofit to z/OS 1.13)
 - DB2 10 APAR PM85944
- CPU reduction for Media Manager I/O (DB2 9 and above)
- Ability to execute code loaded into 1M pageable frames (DB2 11 only and FlashExpress configured)
 - SVL measured 1.8% CPU reduction for an OLTP workload

Shameless Self Promotion



QR code points directly to my blog

Willie Favero

DB2 SME

Data Warehousing for System z Swat Team

IBM Silicon Valley Laboratory

My DB2 Blog

www.it.toolbox.com/blogs/db2zos/

<http://www.WillieFavero.com>



