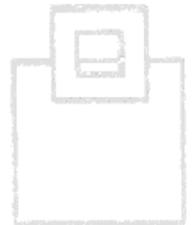
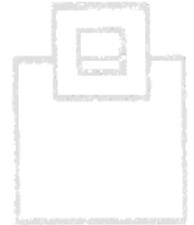
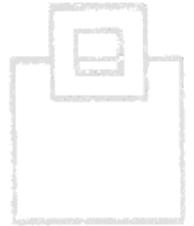


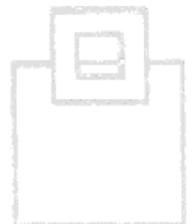
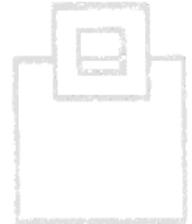
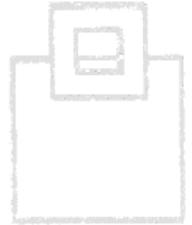
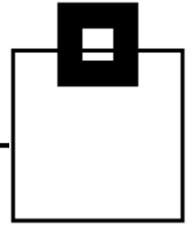
Db2 for z/OS: Lies, Damn lies and Statistics



Ulf Heinrich
SEGUS & SOFTWARE ENGINEERING GmbH

Agenda

- Db2 RUNSTATS basics
- IBM recommendation through the ages
- Db2 RUNSTATS advanced
 - SYSCOLDIST explained
 - RUNSTATS data sampling
 - RUNSTATS PROFILES
- RUNSTATS Rescue



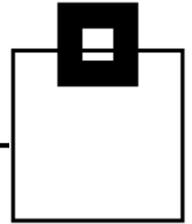
What the optimizer needs to make a choice

1st DDL:

```
SELECT ...  
  FROM ...  
  WHERE ...
```

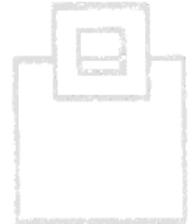
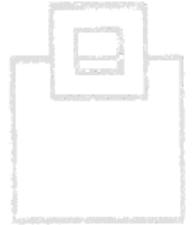
- Which objects are referenced in the SQL
 - SELECT <columns> FROM <table> or <view> or ...
 - WHERE <local predicates>
 - ORDER BY or GROUP BY or UNION or ... <columns>
- Which objects are defined and how
 - INDEX
 - PARTITIONING

What the optimizer needs to make a choice



2nd STATISTICS:

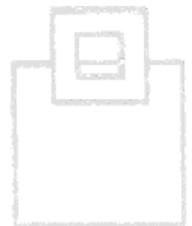
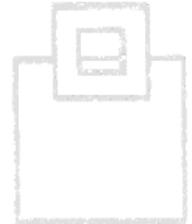
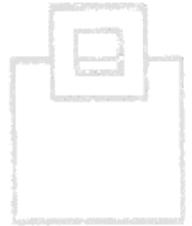
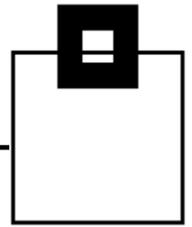
- SYSIBM.SYSCOLDIST
 - SYSIBM.SYSCOLSTATS °
 - SYSIBM.SYSCOLUMNS
 - SYSIBM.SYSINDEXES
 - SYSIBM.SYSINDEXPART
 - SYSIBM.SYSKEYTARGETS (same as SYSCOLUMNS)
 - SYSIBM.SYSKEYTGTDIST (same as SYSCOLDIST)
 - SYSIBM.SYSROUTINES
 - SYSIBM.SYSTABLES
 - SYSIBM.SYSTABLESPACE
 - SYSIBM.SYSTABSTATS
 - SYSIBM.TABLESPACESTATS °
 - SYSIBM.INDEXSPACESTATS °
- degree of parallelism only and, after APAR PK62804, also „sometimes“ used to bound filter factor estimates
 - ° only for dynamic SQL – so far



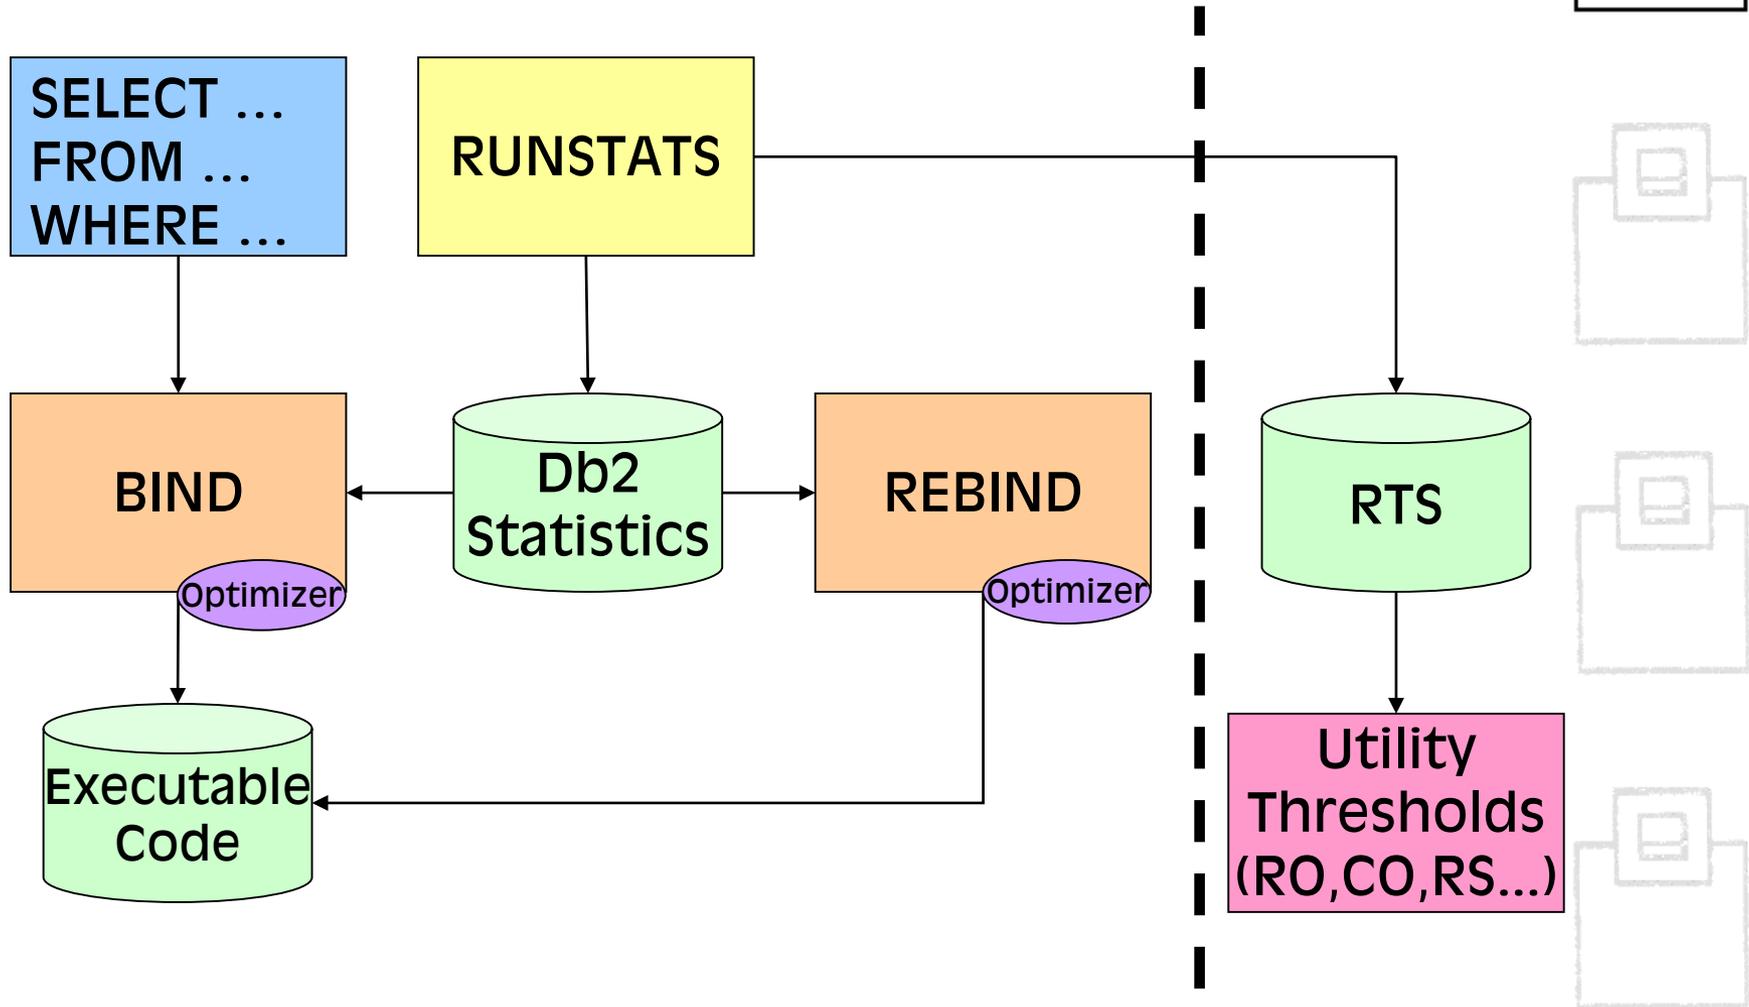
What the optimizer needs to make a choice

3rd ENVIRONMENT:

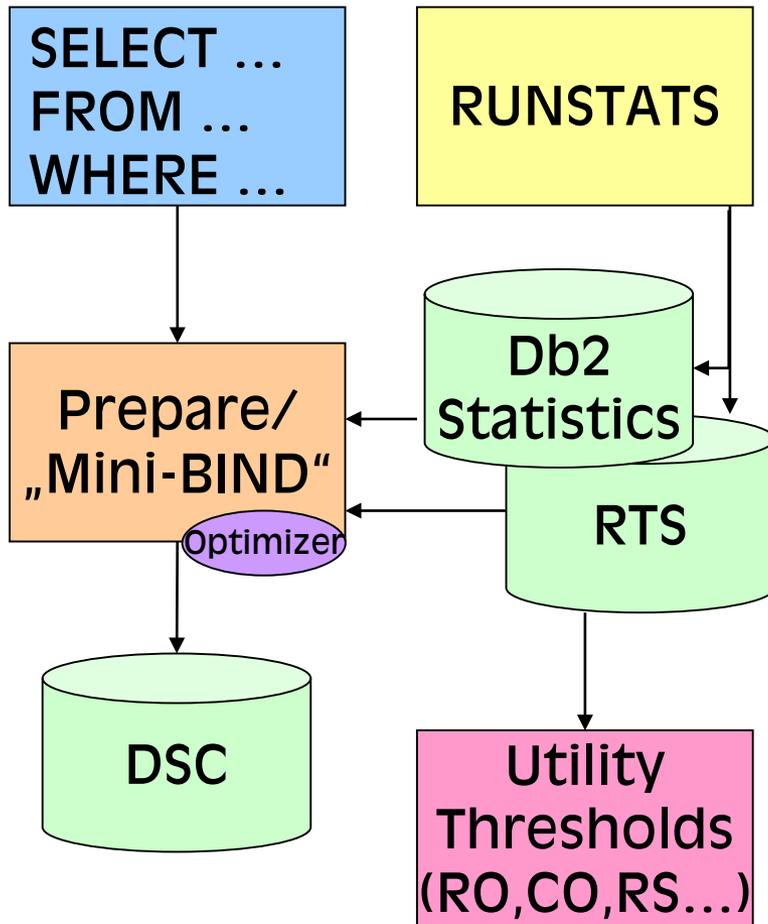
- CP speed
- # of CPUs
- BPs
- RID pool
- Sort pool



Db2 Statistics from an Optimizer's perspective (static SQL)

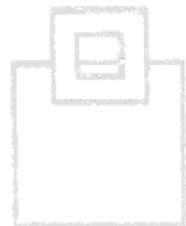
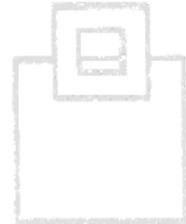
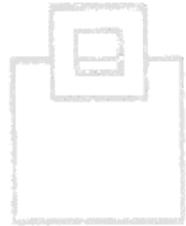
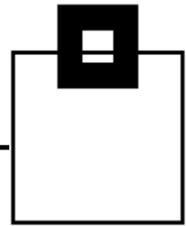


Db2 Statistics from an Optimizer's perspective (dyn. SQL)



Access Paths for dynamic SQL are determined on the fly and stored in the Dynamic Statement Cache (DSC).

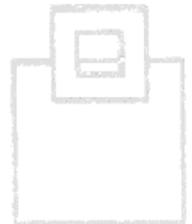
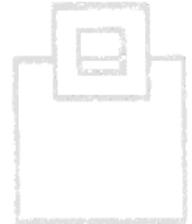
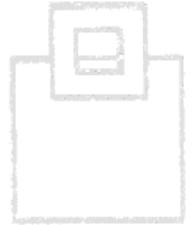
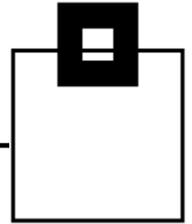
Last Recent Used queue (LRU), RUNSTATS, ALTER, DROP, REVOKE, Db2 RESTART invalidates and flushes the DSC for an object.



Basic RUNSTATS knowledge

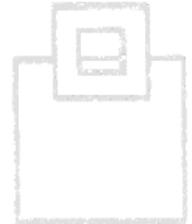
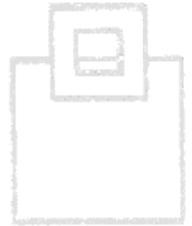
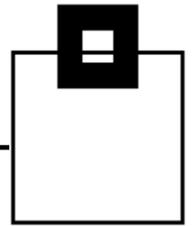
- The RUNSTATS utility
 - Gathers summary information about the characteristics of data in table spaces, indexes and partitions
 - Invalidates the dynamic statement cache*
 - Optionally:
 - Reports the statistics
 - Updates the Db2 catalog
 - Updates the Db2 catalog history tables
 - Updates the Db2 catalog RTS tables

* Default until Db2 12

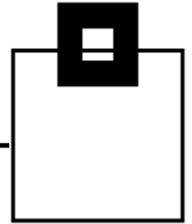


Basic RUNSTATS knowledge

- Different types of RUNSTATS
 - RUNSTATS tablespace
 - RUNSTATS index
 - REORG / LOAD with Inline RUNSTATS
- Different types of statistics
 - Pure access path statistics
 - Those used by BIND in its process of optimization to determine access path
 - Parallel access path statistics
 - Those used by BIND in its process of optimization to determine the degree of parallelism
 - Space statistics
 - Those used by the DBA to monitor space usage; to assist in capacity planning; to determine frequency of REORG; etc.

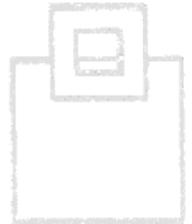


Basic RUNSTATS knowledge

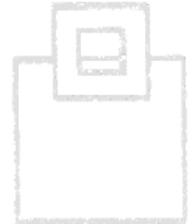


The RUNSTATS Utility

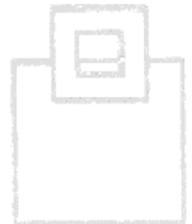
```
RUNSTATS TABLESPACE <DB>.<TS>  
TABLE(<CR>.<TB>)  
  COLGROUP (<CO_A>, <CO_B>)  
  FREQVAL COUNT 10 MOST  
  HISTOGRAM NUMQUANTILES 100  
TABLESAMPLE SYSTEM/SAMPLE 25  
USE PROFILE
```



```
INDEX(ALL)  
KEYCARD FREQVAL NUMCOLS 1 COUNT 10  
          FREQVAL NUMCOLS 2 COUNT 10  
          FREQVAL NUMCOLS 3 COUNT 10  
          HISTOGRAM NUMCOLS 4 NUMQUANTILES 100  
USE PROFILE
```

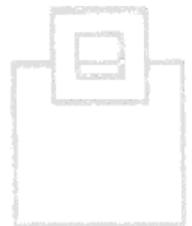
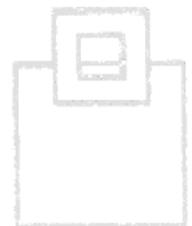
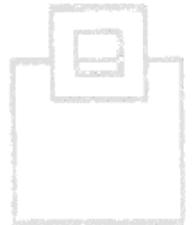
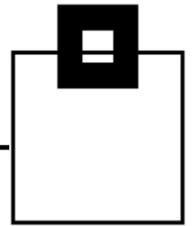


```
REPORT NO  
UPDATE ALL  
HISTORY NONE  
SET/UPDATE PROFILE  
INVALIDATECACHE YES
```



Basic RUNSTATS knowledge

- When to run RUNSTATS
 - After a table is loaded
 - After an index is physically created
 - After running extensive updates, deletions, or insertions in a table space
 - After running any of the following utilities without collecting inline statistics: RECOVER TABLESPACE, REBUILD INDEX, or REORG TS/IX
 - Before running REORG with the OFFPOSLIMIT, INDREFLIMIT, or LEAFDISTLIMIT options
 - After running the ALTER TABLE ROTATE PARTITION statement run RUNSTATS with REORG
 - When the distribution of the data changes
 - When the values over which the data is distributed change
 - When RUNSTATS is enhanced (clusterratio → DRF)



The importance of RUNSTATS

- If statistics are old, inconsistent, or incomplete, the optimizer can't make a proper decision
 - If no frequency data exists in the SYSCOLDIST for example, Db2 would assume that all values are equally distributed
 - Default filter factor of 1/5 (1/COLCARD) or 20%
 - ADULT under estimated by 50%
 - INFANT over estimate by 400%

Category	Frequency
Infant	5%
Child	15%
Adolescent	25%
Adult	40%
Senior	15%

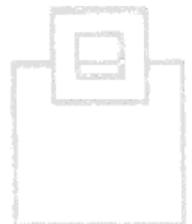
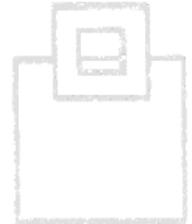
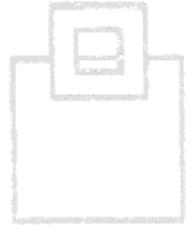
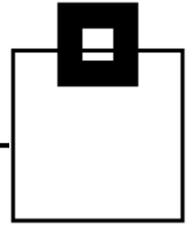


OBJECTS IN MIRROR ARE
LARGER THAN THEY APPEAR

The importance of RUNSTATS AND REBIND!

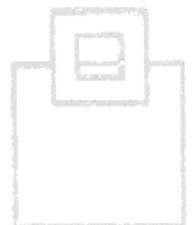
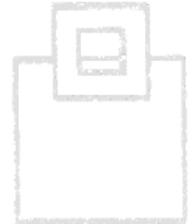
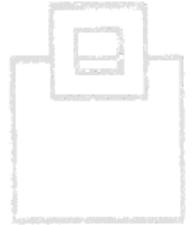
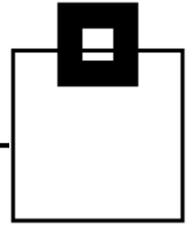
- Unless the entire workload is dynamic, RUNSTATS only updates the optimizer base, but not the access path
 - REBIND is required to
 - exploit statistics updates
 - exploit optimizer enhancements
 - enable new Db2 features and enhancement
 - Some shops are still frightened about REBINDs, because of
 - access path degradations
 - difficulties in recovering prior access paths
 - differences in managing dynamic/static SQL
 - changed behavior of Db2 functions (e.g. BiFs)

But various Db2 options can help...

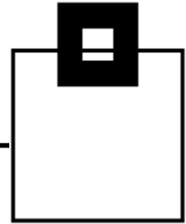


Db2 features to understand and manage access paths

- We need to be able to pre-screen access path changes resulting from statistics and code changes
 - Since Db2 9 we can use PLAN MANAGEMENT to fallback to prior runtime structures
 - Since Db2 10
 - we can use AP recycling and comparison to manage and to be aware of access path changes
 - we have BIND QUERY to add hints to dynamic queries
 - Since Db2 12 we have stabilized dynamic queries to freeze runtime structures
 - Since Db2 collections we have the opportunity to REBIND (distributed) packages into a separate collection to compare and switch

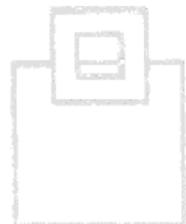
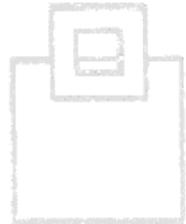


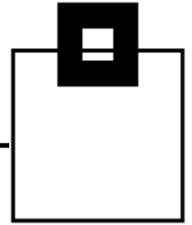
Columns used for access path decisions



<u>SYSCOLDIST /</u> <u>SYSKEYTGTDIST</u>	<u>SYSCOLSTATS</u>	<u>SYSTABLES</u>
CARDF	COLCARD	CARDF
COLGROUPCOLNO / KEYGROUPKEYNO	HIGHKEY	EDPROC*
COLVALUE / KEYVALUE	LOWKEY	NPAGES
FREQUENCYF	<u>SYSINDEXES</u>	NPAGESF
HIGHVALUE	<u>CLUSTERING*</u>	PCTROWCOMP
LOWVALUE	CLUSTERRATIO	<u>SYSTABLESPACE</u>
NUMCOLUMNS /	CLUSTERRATIOF	NACTIVE
NUMKEYS	DATAREPEATFACTORF	NACTIVEF
QUANTILENO	FIRSTKEYCARDF	<u>SYSTABSTATS</u>
STATSTIME	FULLKEYCARDF	CARD
<u>SYSCOLUMNS /</u>	NLEAF	CARDF
<u>SYSKEYTARGETS</u>	NLEVELS	NPAGES
COLCARDF / CARDF	<u>SYSROUTINES</u>	<u>SYSTABLESPACESTATS</u>
HIGH2KEY	CARDINALITY*	<u>SYSINDEXSPACESTATS</u>
LOW2KEY	INITIAL_INSTS*	TOTALENTRIES
n/a / STATS_FORMAT	INITIAL_IOS*	TOTALROWS
<u>SYSINDEXPART</u>	INSTS_PER_INVOC*	
<u>LIMITKEY*</u>	IOS_PER_INVOC*	

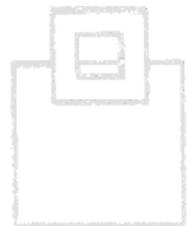
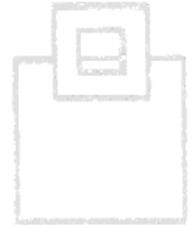
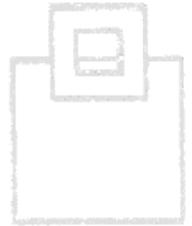
* Columns are not updated by RS
 _ Columns are not updatable





Correlations in the catalog (DB2 Administration Guide)

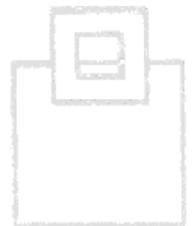
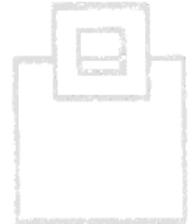
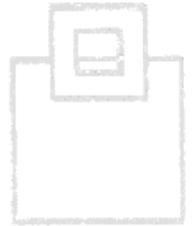
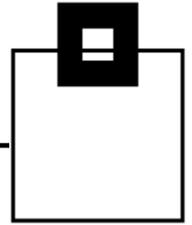
- Relationships exist among certain columns of certain tables:
 - Columns within SYSCOLUMNNS
 - Columns in the tables SYSCOLUMNNS and SYSINDEXES
 - Columns in the tables SYSCOLUMNNS and SYSCOLDIST (wellactually SYSFIELDS in those days!)
- If you plan to update some values, keep in mind the following correlations:
 - COLCARDF and FIRSTKEYCARDF/FULLKEYCARDF
 - COLCARDF, LOW2KEY and HIGH2KEY. For non-default COLCARDF



IBM recommendations Db2 V3.1

Correlations in the catalog (DB2 Administration Guide)

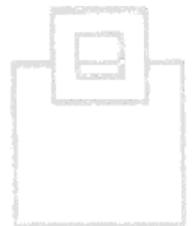
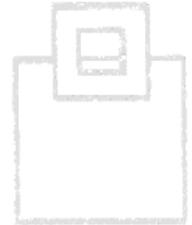
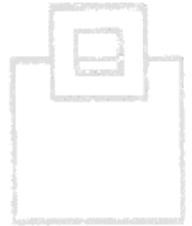
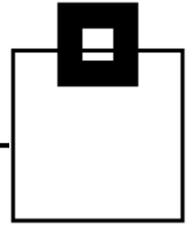
- No change

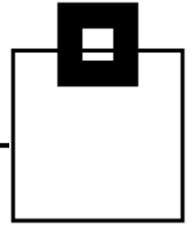


IBM recommendations Db2 V4.1

Correlations in the catalog (DB2 Administration Guide)

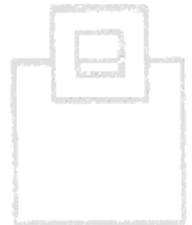
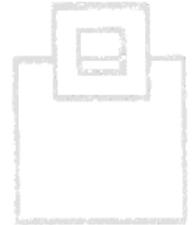
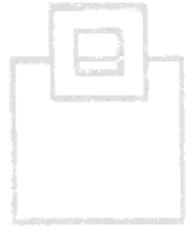
- No change





Correlations in the catalog (DB2 Administration Guide)

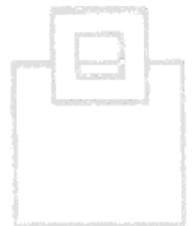
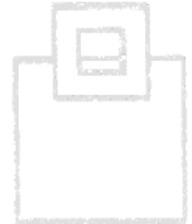
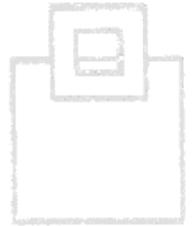
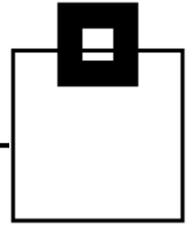
- Relationships exist among certain columns of certain tables:
 - Columns in the tables SYSCOLUMNS, SYSCOLDIST, and SYSINDEXES
- If you plan to update values, keep in mind the following correlation:
 - CARDF in SYSCOLDIST. CARDF is related to COLCARDF and FIRSTKEYCARDF and FULLKEYCARDF. It must be at minimum:
 - A value between FIRSTKEYCARDF and FULLKEYCARDF if the index contains the same set of columns
 - A value between $\text{MAX}(\text{colcardf of each col})$ and the product of all the columns COLCARDFs in the group



IBM recommendations Db2 V6.1

Correlations in the catalog (DB2 Administration Guide)

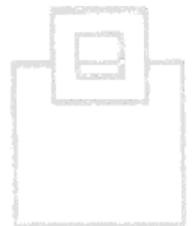
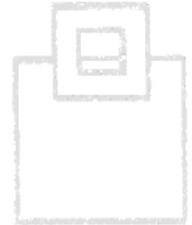
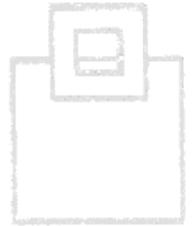
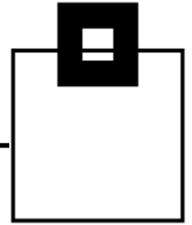
- No change

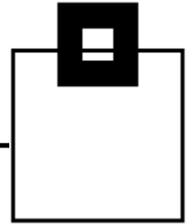


IBM recommendations Db2 V7.1

Correlations in the catalog (DB2 Administration Guide)

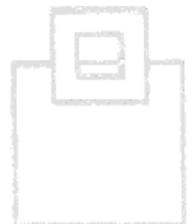
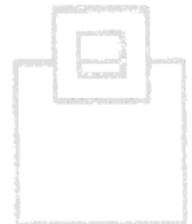
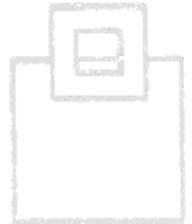
- No change

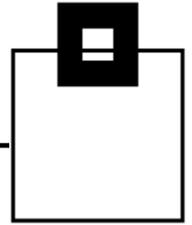




Correlations in the catalog (DB2 Administration Guide)

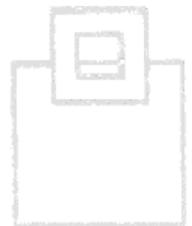
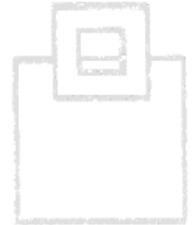
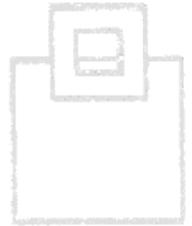
- If you plan to update values, keep in mind the following correlations:
 - The COLCARDF, LOW2KEY, and HIGH2KEY gained a note: „If the COLCARDF is 1 or 2 DB2 uses LOW2KEY and HIGH2KEY as domain statistics to generate frequencies“
 - CARDF in SYSTABLES. CARDF must be equal or larger than any other cardinalities, such as COLCARDF, FIRSTKEYCARDF, FULLKEYCARDF, and CARDF in SYSCOLDIST
 - FREQUENCYF and COLCARDF or CARDF. The number of frequencies collected must be less than or equal to COLCARDF for the column or CARDF for the column group
 - FREQUENCYF. The sum of frequencies collected for a column or column group must be less than or equal to 1





In the „Filter factors and catalog statistics“ chapter (DB2 Administration Guide)

- Recommendation: If query performance is not satisfactory, consider the following actions:
 - Collect cardinality statistics on all columns that are used as predicates in a WHERE clause.
 - Collect frequencies for all columns with a low cardinality that are used as COL op literal predicates.
 - Collect frequencies for a column when the column can contain default data, the default data is skewed, and the column is used as a COL op literal predicate.
 - Collect KEYCARD on all candidate indexes.
 - Collect column group statistics on all join columns.



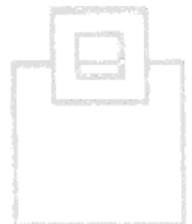
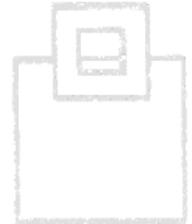
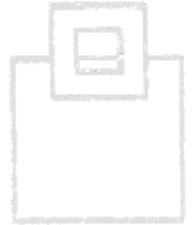
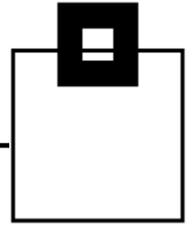
IBM recommendations Db2 9

Correlations in the catalog in Chapter 36 of the new book: „DB2 Performance Monitoring and Tuning Guide“

- No change

Filter factors and catalog statistics in Chapter 33 of the new: „DB2 Performance Monitoring and Tuning Guide“

- New section all about HISTOGRAM statistics.

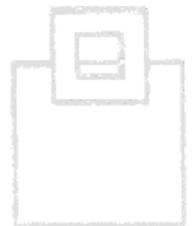
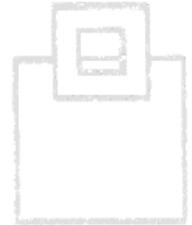
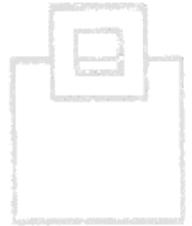
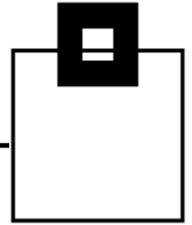


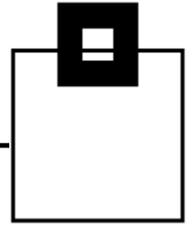
Correlations in the catalog in chapter 37 of: „Managing Performance“

- No change

Filter factors and catalog statistics in chapter 34 of: „Managing Performance“

- KEYCARD use deleted. Now it is always active.



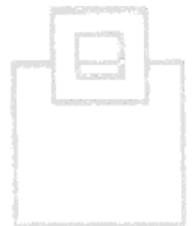
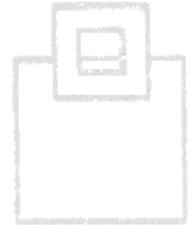
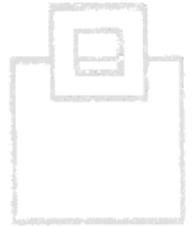


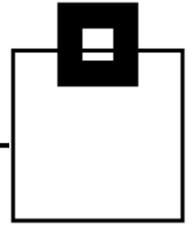
Correlations in the catalog in chapter 37 of: „Managing Performance“

- No change

Filter factors and catalog statistics in chapter 34 of: „Managing Performance“

- SYSSTATFEEDBACK usage introduced for missing or conflicting statistics.



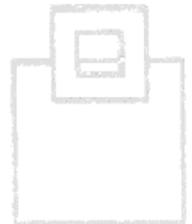
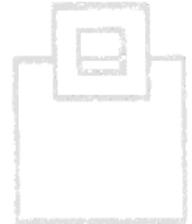
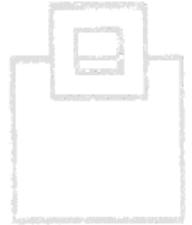


Correlations in the catalog in chapter 37 of: „Managing Performance“

- No change

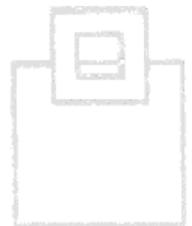
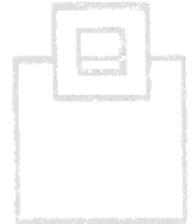
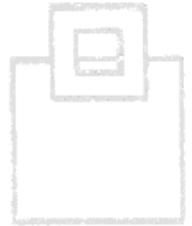
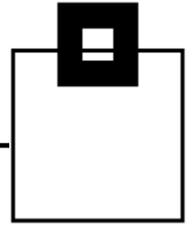
Filter factors and catalog statistics in chapter 34 of: „Managing Performance“

- SYSSTATFEEDBACK got one new reason - „STALE“.



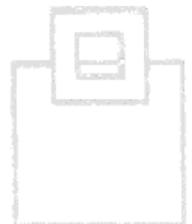
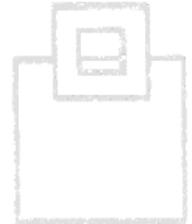
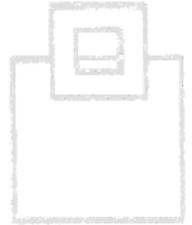
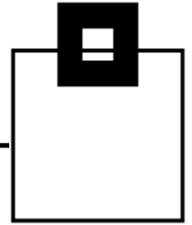
SYSCOLDIST contents explained

- SYSCOLDIST is used for:
 - Frequencies (good for COL op literal)
 - Cardinalities (good for everything!)
 - Histograms (good for COL op literal)
- Column TYPE can contain:
 - „C“ for cardinality values
 - „F“ for frequency values
 - „N“ for non-padded frequency values
 - „H“ for histogram values
- Regardless of the TYPE value, columns TBOWNER, TBNAME, NAME, COLGROUPOCOLNO, NUMCOLUMNS, and STATSTIME are used for the same purpose.



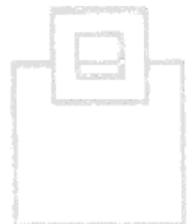
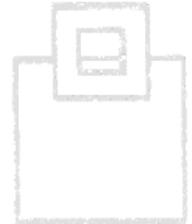
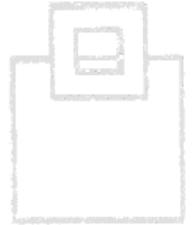
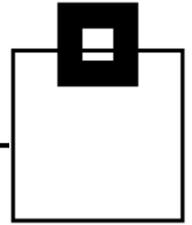
SYSCOLDIST contents explained

- **TBOWNER, TBNAME, and NAME** (first column name only) columns are also the non-unique index.
- **COLGROUPOCOLNO** for a single column object is an empty string, for a multi-column object contains a string of two byte **SMALLINT**, which contains the column numbers from the original table.
- **NUMCOLUMNS** is the number of columns in this group.
- **STATSTIME** is the time when **RUNSTATS** inserted this entry, which is also used when there is a complete duplicate so that Db2 uses the last inserted value.



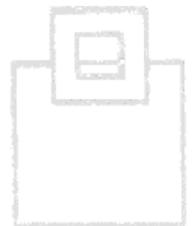
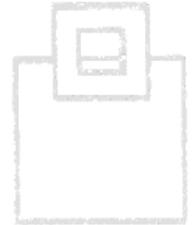
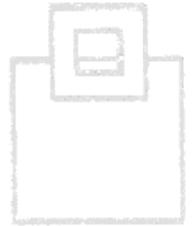
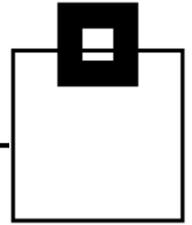
SYSCOLDIST contents explained

- For FREQUENCY data:
 - COLVALUE contains the actual data from the 1 – n columns in the group simply concatenated together.
 - This data might be readable or it might not.
 - Caution must be used when inserting or changing this data as numeric and date, time forms must be the internal Db2 format and not the external format
 - For example, the high bit must be flipped for numbers, and watch out for NULLable columns!
 - FREQUENCYF contains a floating point value between 0.0 and 1.0 which is the frequency that this value occurs for this group in the table.



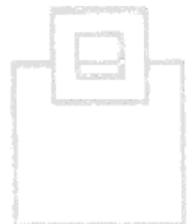
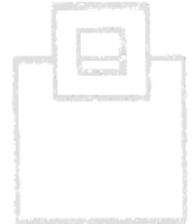
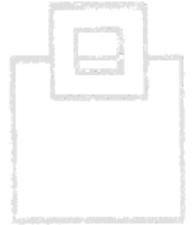
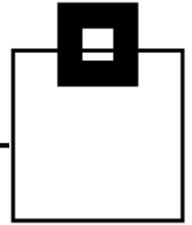
SYSCOLDIST contents explained

- For CARDINALITY data:
 - CARDF contains a floating point value
 - Which is how many different values for this group occur in the data.



SYSCOLDIST contents explained

- For HISTOGRAM data:
 - COLVALUE contains the actual data from the 1 – n columns in the quantile simply concatenated together.
 - This data might be readable or it might not.
 - Caution must be used when inserting or changing this data as numeric and date, time forms must be the internal Db2 format and not the external format
 - For example, the high bit must be flipped for numbers, and watch out for NULLable columns!
 - FREQUENCYF contains a floating point value between 0.0 and 1.0 which is the frequency that this value occurs for this quantile.
 - HIGHVALUE and LOWVALUE contain the upper and lower boundaries of this QUANTILENO.
 - QUANTILENO is the ordinary number of this quantile (1 - 100)



RUNSTATS data sampling

- Use of SAMPLE
 - Only for non-indexed columns
 - Test results show it has an impact on indexed columns as well.
 - Problem with SAMPLE: Even using SAMPLE 100, Db2 still gets it wrong...

Col	1	5	10	20	21	22	23	24	25	99	100
TBNAME	62	121	162	178	178	180	180	178	180	180	180
CLNAME	11647	11647	5442	4854	11647	5057	5340	3293	6707	7210	7168

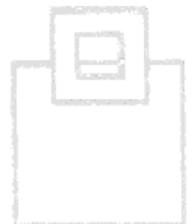
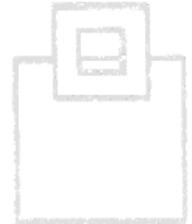
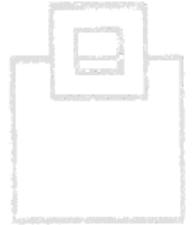
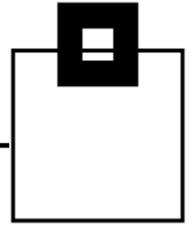
- The actual values are 181 and 6871 – be careful using SAMPLE

RUNSTATS data sampling

- Use of TABLESAMPLE SYSTEM
 - TABLESAMPLE SYSTEM AUTO recommended for any table space with over 500,000 pages but not for LOBs.
 - The following is from the Utility Guide:

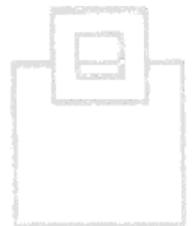
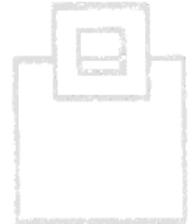
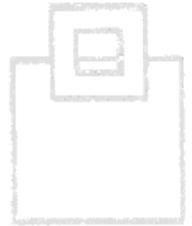
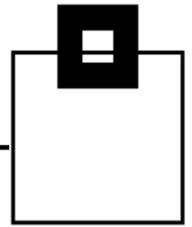
„When RUNSTATS TABLESPACE TABLE is run on a multi-table table space or a table space that is segmented and not partitioned, the TABLESAMPLE keyword is not used. Instead, all pages are scanned to collect statistics.“

Clear? Simply put – UTS only!



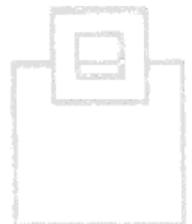
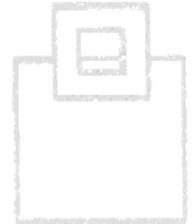
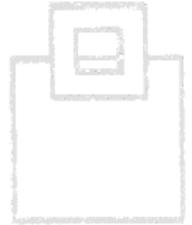
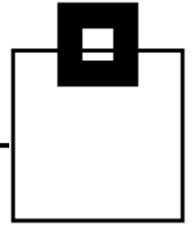
RUNSTATS PROFILES

- Set up your profiles first, then you can update/delete them making the RUNSTATS syntax easier.
- If you have enabled the auto-update of profiles, which is **on** by default, then you must make sure that your profiles are not growing out of control!
- All profile data stored in `SYSIBM.SYSTABLES_PROFILES`.
- From then on, just adding `USE PROFILE` with/without the `INCLUDE(NPI)` syntax, and no other columns, colgroups, freqs etc., is all you need.
- If no profile is found you will get an automatic
COLUMN ALL INDEX ALL
style runstats.



RUNSTATS PROFILES

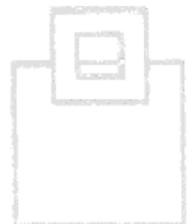
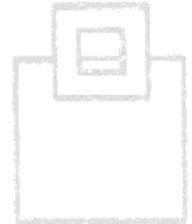
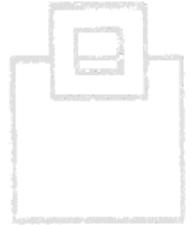
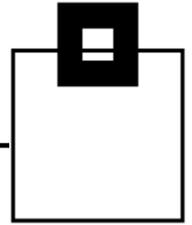
- Applying statistics recommendations to statistics profiles automatically
 - This is new in Db2 12. To enable this feature you do nothing as it is, by default on. This is, in my opinion, a bad thing!
 - ZPARM STATFDBK_SCOPE set to ALL by default
ZPARM STATFDBK_PROFILE set to YES by default
SYSIBM.SYSTABLES column STATS_FEEDBACK set to Y by default
 - Out-of-the box it starts automatically creating (for TYPE='C' with NUMCOLS > 1 and TYPE='F' or 'H') and updating profiles...



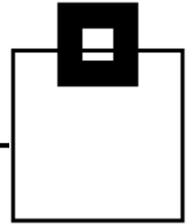
RUNSTATS Rescue

- RUNSTATS became definitely more powerful, but also more complex
- There is no common approach in managing dynamic AND static access plans
- If critical access plans degrade in production we need a quick and solid approach to recover

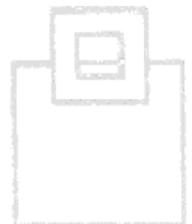
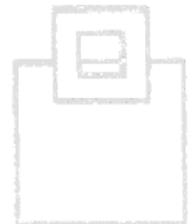
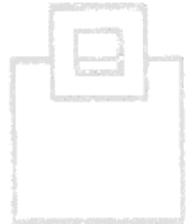
→ Consider RUNSTATS Rescue



RUNSTATS Rescue

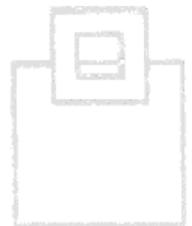
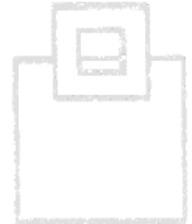
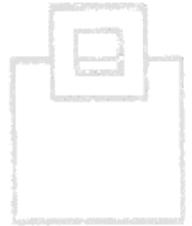
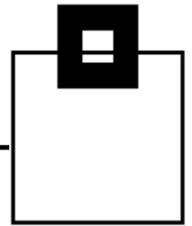


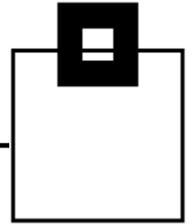
- RUNSTATS Rescue concept:
 - Every day, copy the complete Production Db2 Catalog Optimizer relevant data to a Generation Dataset (refer to slide 15)
- If a RUNSTATS leads to a bad access plan:
 - Locate in this data the required catalog info to restore the Db2 Catalog columns to a point-in-time **before** a RUNSTATS killed an access plan. Consider using EXPLAIN or SYSPACKDEP to limit the search.
 - Flush the bad guy from the DSC and/or REBIND the affected packages.
 - Prior access plan is regenerated and reactivated



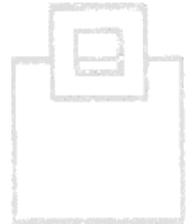
Conclusion

- Be careful with your Db2 Catalog statistics
- Garbage In - Garbage Out!
- Guard your access paths as if they are good!
- Control the number and complexity of COLGROUPS – It impacts PREPARE time
- Plan a way to get your old statistics back as quickly as you can!

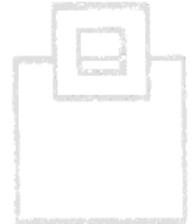




**"Lies, damn lies, and statistics" – Benjamin Disraeli,
Prime Minister of England (1868, 1874-1880)**



The above line may, or may not, have been spoken well over 100 years ago, but the need for statistics and, above all else, accurate statistics is more important than ever in the Db2 world of today.



Ulf Heinrich
u.heinrich@segus.com

