

Under the hood of dynamic and adaptive database features



Stefan Koehler

About me

Stefan Koehler

- Independent Oracle performance consultant and researcher
- 15+ years using Oracle RDBMS - Independent since 2011
- Oracle performance and internals geek
- Main interests: Cost based optimizer and Oracle RDBMS internals



OakTable.net



Services: “All about performance & troubleshooting”

- Oracle performance tuning (e.g. Application, CBO, Database, Design, SQL)
- Oracle core internals researching (e.g. DTrace, GDB, Perf, etc.)
- Troubleshooting nontrivial Oracle RDBMS issues (e.g. Heap dumps, System state dumps, etc.)
- Services are mainly based on short-term contracting



www.soccs.de



contact@soccs.de



[@OracleSK](https://twitter.com/OracleSK)

Agenda

- Cardinality Feedback (CFB) / Statistics Feedback
 - Basics
 - Evolution
 - Tracing for troubleshooting and researching purpose
 - Known limitations including work around
- Dynamic Sampling (DS)
 - Why dynamic sampling
 - New dynamic sampling code (implementation) and its side effects



Disclaimer: Almost everything is based on research and testing. Test it yourself - with your release and operating system – always! Do not trust anybody! 😊

Cardinality / Statistics Feedback - Basics

- Cardinality Feedback aims to address cardinality misestimates by the Oracle CBO due to complex predicates, and inaccurate or missing statistics.
- During optimization, certain types of estimates that are known to be low quality are noted, and monitoring is enabled for this child cursor - at the end of (first) execution of this cursor, the estimates are compared to the actual cardinalities.
- If an estimate is found to be significantly different from the actual value, then the correct value is stored for later use. On subsequent executions, the query is re-optimized, and the correct cardinality is substituted for the usual estimate.
- If the estimate is found to be relatively close to the original estimate after the first execution, then the monitoring is disabled for future executions.



Cardinality / Statistics Feedback - Evolution

Automatic Re-optimization

Complex Predicates?

What is significantly? 5 or 10 times?

During the first execution of a SQL statement, an execution plan is generated as usual. During optimization, certain types of estimates that are known to be of low quality (for example, estimates for tables which lack statistics or tables with complex predicates) are noted, and monitoring is enabled for the cursor that is produced. If feedback monitoring is enabled for a cursor by the system, cardinality estimates in the plan are compared to the actual cardinalities seen during execution. If estimates are found to differ significantly from the actual cardinalities, then the optimizer looks for a replacement plan on the next execution. The optimizer will use the information gathered during the previous execution to help determine an alternative plan. The optimizer can re-optimize a query several times, each time learning more and further improving the plan. Oracle Database 12c supports multiple forms of re-optimization.



[Oracle White Paper "Optimizer with Oracle Database 12c Release 2" - June 2017](#)

ORACLE[®] 11^g
DATABASE
11.2.0.1 / (11.1.0.7)



ORACLE[®] 12^c
DATABASE
12.1.0.1



ORACLE[®] 12^c
DATABASE
12.2.0.1

- Introduced as Cardinality Feedback
- Single Table CFB only
- Feedback lost after Shared Pool Flush

- Renamed to Statistics Feedback
- Single Table & Join Statistics Feedback
- Feedback is remembered as SPD (object level) and used by default

- Single Table & Join Statistics Feedback
- Feedback is remembered as SPD (object level) and not used by default (except ADS)

Cardinality / Statistics Feedback - Tracing



```
shell > oerr ora 10507  
10507, 00000, "Trace bind equivalence logic"  
// *Cause:  
// *Action:
```

Cardinality / Statistics Feedback - Limits

1. Execute SQL through OCI
2. Execute SQL through PL/SQL engine
3. Execute “small” SQL queries (≥ 12.2)




Dynamic Sampling (DS) - Why?

- Dynamic sampling (DS) was introduced to improve the optimizer's ability to generate good execution plans - it is used when regular statistics are not sufficient to get good quality cardinality estimates.
- The optimizer decides whether to use DS or not during the compilation of a SQL.



LEVEL	WHEN DYNAMIC STATISTICS WILL BE USED	SAMPLE SIZE (BLOCKS)
0	Switches off dynamic statistics	N/A
1	At least one non-partitioned table in the statement has no statistics	32
2 (default)	One or more tables in the statement have no statistics This setting is suitable for the majority of systems	64
3	Any statement that meets level 2 criteria and any statement that has one or more expressions used in the where clause predicates e.g. Where substr(CUSTLASTNAME,1,3) or Where a + b =5	64
4	Any statement that meets level 3 criteria and any statement that has complex predicates. An OR or AND operator between multiple predicates on the same table	64
5	Any statement that meets level 4 criteria	128
6	Any statement that meets level 4 criteria	256
7	Any statement that meets level 4 criteria	512
8	Any statement that meets level 4 criteria	1024
9	Any statement that meets level 4 criteria	4086
10	All statements	All Blocks
11	The database determines automatically if dynamic statistics are required	Automatically determined

Dynamic Sampling (DS) - New code

- New implementation is called adaptive dynamic sampling (ADS) and is available since 12c but also back-ported to 11.2.0.4. It is enabled with `OPTIMIZER_DYNAMIC_SAMPLING = 11` and may run several sampling queries (Complexity & Indexes). 
- It is not limited to single-table cardinality estimations anymore (join cardinalities and query block cardinalities).
- In case of PX plan the optimizer can automatically increase the dynamic sampling level to level 11 (from default level 2) since 12.1.0.2. Before 12.1.0.2 it was between 2 and 11.
- SQL Plan Directives use the new ADS code since 12.1.0.2.
- The dynamic sampling phase can't exceed a given time limit and it depends on whether the SQL statement is in the cursor cache or AWR. Each sampling query can't exceed a given time slice (ORA-10173).

Questions and answers



The slides including the SQL demo scripts can be found on website
www.soccs.de/public/talk/



www.soccs.de



contact@soccs.de



[@OracleSK](https://twitter.com/OracleSK)

References

- <https://blogs.oracle.com/optimizer/cardinality-feedback>
- https://antognini.ch/2011/09/impact-of-statistics_level-on-cardinality-feedback-and-adaptive-cursor-sharing/
- <https://antognini.ch/2016/10/adaptive-query-optimization-configuration-parameters-preferences-and-fix-controls/>
- <https://blog.dbi-services.com/when-automatic-reoptimization-plan-is-less-efficient/>
- <https://blog.dbi-services.com/12cr2-no-cardinality-feedback-for-small-queries/>
- <http://oracle-randolf.blogspot.de/2017/01/12c-adaptive-joins-plus-statistics.html>
- <https://blogs.oracle.com/optimizer/dynamic-sampling-and-its-impact-on-the-optimizer>
- <https://blogs.oracle.com/datamining/to-sample-or-not-to-sample-part-2>
- <https://jonathanlewis.wordpress.com/2010/03/12/sample-clause/>
- <https://jonathanlewis.wordpress.com/2010/02/23/dynamic-sampling/>
- http://www.itoug.it/wp-content/uploads/2016/05/AdaptiveDynamicSampling_OTNMI_2016.pdf
- <https://mauro-pagano.com/2016/11/28/something-new-about-sql-plan-directives-and-12-2/>