

A large, stylized blue chevron graphic pointing to the right, serving as a background for the main title text.

MANAGING COST-BASED OPTIMIZER STATISTICS FOR PEOPLESOFT

**DAVID KURTZ
UKOUG PEOPLESOFT ROADSHOW 2018**

WHO AM I

Accenture Enkitech Group

Performance tuning

- PeopleSoft ERP
- Oracle RDBMS

Book

- www.go-faster.co.uk
- blog.psftdba.com

 **Oak Table**

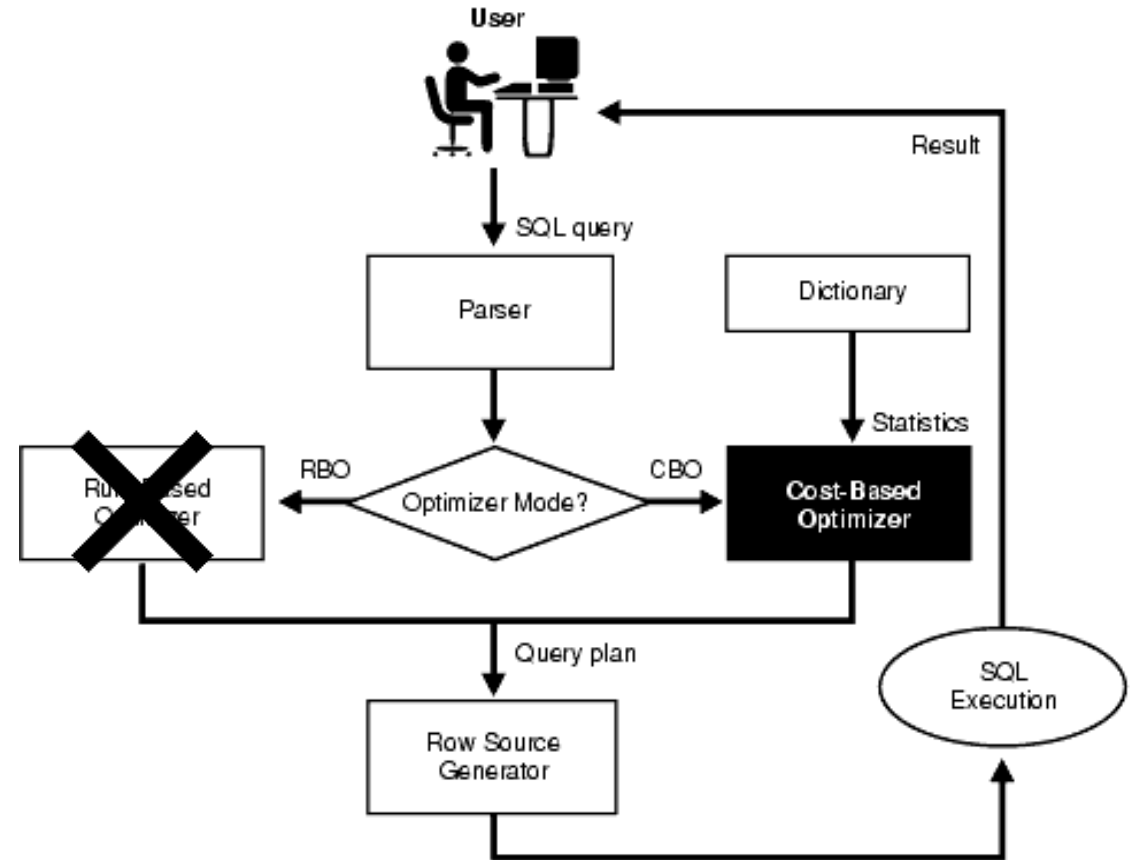


AGENDA

- How the Oracle Optimizer works
 - What is Cost
 - What are statistics
- Collecting Stats
 - Maintenance Window
 - At runtime
 - My recommendations
- Management Packages
 - GFCPSSTATS11
 - Oracle's PSCBO_STATS
- Not Statistics (but related)
 - Adaptive Optimization
 - SQL Plan Directives
 - Parameters

SQL PROCESSING


- The Parser checks both syntax and semantic analysis.
- The Optimizer uses costing methods to determine the most efficient way of producing the result of the query.
- The Row Source Generator receives the optimal plan from the optimizer and outputs the execution plan for the SQL statement.
- The SQL Execution Engine operates on the execution plan associated with a SQL statement and then produces the results of the query.
- (Diagram from [9i Tuning Guide](#) – things are more complex today)



COST-BASED OPTIMIZATION

Car GPS

- Calculate route A → B
- It calculates all possible routes
- Chooses the least expensive
- Each arc in the route has a cost
- More expensive routes are pruned out as early as possible
- ([Dijkstra's Algorithm](#) – diagram from Wikipedia)



A → B : 1 → 3 → 6 → 5

WHAT DO WE MEAN BY 'COST'?

Car GPS

- Distance
- Time
- Blended measure, also including
 - Fuel Cost
 - Tolls
- Some route planners also factor in time based congestion.

SQL Optimizer

- Time
 - Physical I/O
 - CPU Cost
 - Built in heuristics
- Very roughly, optimizer cost is time, where the unit of time is the duration of a single block read

Statistics

- Number of rows in a table
- Number of blocks in a table
- Range of data values on a column (min/max)
- Distribution of data values on a column
- Number of rows in an index
- Number of distinct keys
- How many table blocks will we visit if we scan through an index (clustering factor).

COST-BASED OPTIMIZER STATISTICS

We need object statistics

They need to be accurate enough for the optimizer to make good decisions

But data changes, so the statistics need to reflect those changes, often so plans remain the same

“The only time it is safe to collect statistics is when to do so would make no difference”

- (sometimes known as the Ensor Paradox)

Oracle Supplied PL/SQL Packaged Procedure *DBMS_STATS* that calculates and manages statistics.

- Calculate Statistics on
 - Object
 - Objects in schema
 - Whole database in schema
 - Eg. *DBMS_STATS.GATHER_TABLE_STATS()*
- Delete Statistics
- Lock Statistics
- Export/Import Statistics

MAINTAINING OPTIMIZER STATISTICS

Up to **Oracle 9i** it was common to write custom scripts to refresh all statistics

- Could consume a lot of resource and take a long time.
- Inevitably computed statistics on too many tables too often.

Oracle 10g/11g

- Maintenance Window
 - 10pm-6am weekdays, all day at weekends
- Stale (if 10% change)
 - From 11g can be set per table
- Statistics Locking
 - Tables omitted from schema-wide gather processes if statistics locked
 - use FORCE option to override lock

WHAT STATISTICS TO COLLECT

DBMS_STATS parameters

- *METHOD_OPT* – Histograms
 - Only collect when necessary. Might need to restrict
- *DEGREE* – parallelism
- *GRANULARITY* – partitions

Only via table statistics preferences

- *INCREMENTAL* - partitioning
- *STALE_PERCENT* - Stale Threshold

DO NOT SPECIFY

- *SAMPLE_SIZE*
 - always use default of *DBMS_STATS.AUTO_SAMPLE_SIZE* to invoke new hash-based number of distinct values (NDV) algorithm
- *CASCADE*
 - Controls whether index stats collected at same time as table stats
 - true by default

COLLECTING STATISTICS AT RUNTIME

Temporary working storage tables in all batch processing but especially Application Engine

- Truncated,
- repopulated,
- and then queried.

Explicitly collect stats during batch processes

- In AE with *%UPDATESTATS*
 - Command in DDL Models 4 & 5
 - In some delivered programs
 - You may need to add more
- Some Cobol programs collect statistics
 - From 8.55 also uses DDL Model

General Principle:

Tell the optimizer the truth

When you work with any table:

- Populate the table
- Collect Statistics on the table*
- Use the table

DELIVERED DDLMODEL

Model 4

- %UpdateStats([table],LOW);

```
DBMS_STATS.GATHER_TABLE_STATS  
(ownname=> [DBNAME], tabname=>[TBNAME]  
, estimate_percent=>1  
, method_opt=> 'FOR ALL INDEXED COLUMNS SIZE 1'  
, cascade=>TRUE);
```

Model 5

- %UpdateStats([table],HIGH);

```
DBMS_STATS.GATHER_TABLE_STATS  
(ownname=> [DBNAME], tabname=>[TBNAME]  
, estimate_percent=> dbms_stats.auto_sample_size  
, method_opt=> 'FOR ALL INDEXED COLUMNS SIZE 1'  
, cascade=>TRUE);
```

Problems

- LOW uses 1% sample size
 - this may be greater than *AUTO_SAMPLE_SIZE*
 - It disables the new hash-based number of distinct values algorithm introduced in 11g
 - so still uses a sort, that can have a higher overhead
- *FOR ALL INDEXED COLUMNS SIZE 1*
 - Only updates columns statistics on index columns.
 - Non-indexed columns are not updated and will thus be out of date.
 - Removes all histograms from a table, irrespective of the statistics preferences

RECOMMENDED SIMPLE DDLMODEL

Model 4 & 5

- *%UpdateStats([table],LOW|HIGH);*
- Same for both

```
DBMS_STATS.GATHER_TABLE_STATS
```

```
(ownname=> [DBNAME], tabname=>[TBNAME]
```

```
,force=>TRUE);
```

Recommendations

- No parameters are specified
- Defaults are always used
 - As amended by table statistics preferences
 - *AUTO_SAMPLE_SIZE*
 - *CASCADE=TRUE*
 - *METHOD_OPT=GATHER AUTO*
- Specify *FORCE=>TRUE* to override locked statistics

WHAT HAPPENS IF YOU DON'T COLLECT STATISTICS AT RUNTIME?

Yesterday

- Ran a batch program
- It truncated and populated a table used for temporary working storage
- So the statistics on the table are now stale

Today

- Ran the same batch program again
- It truncated the working storage table
 - This does not remove the statistics
- It repopulated the working storage table
 - New Data and Old Statistics

Last Night

- Statistics collected on during maintenance window
- On tables with unlocked statistics
 - without statistics or
 - with stale statistics
- Statistics updated on working storage table in batch process

Conclusion:

- The only time to collect statistics on working storage tables is in the process that populates them.
- It is too late to collect statistics on temporary working storage tables during the maintenance window

WHAT IS THE PROBLEM WITH STALE STATISTICS ON TEMPORARY TABLES?

For most temporary working storage tables

- PROCESS_INSTANCE is the first column in the unique key.
- All the rows will have the same PROCESS_INSTANCE
- Most processes will select rows from table with a criterion on PROCESS_INSTANCE
 - let's assume there are no other predicates

```
SELECT * FROM <table>
```

```
WHERE process_instance = %PROCESS_INSTANCE
```

- Every column has a min/max value

Statistics on the temporary table

- Min/max values on PROCESS_INSTANCE are yesterday's process
- Therefore current values of PROCESS_INSTANCE in the table > column maximum value in the statistics
- The optimizer will
 - calculate that no rows exist for today's process instance.
 - and assume that it will get a single row.
 - It will probably use an index
- In fact
 - All the rows will be returned.
 - It might have been better to full scan this table

DYNAMIC STATISTICS OPTIMIZER DYNAMIC SAMPLING

Designed to augment rather than replace optimizer statistics

- [Oracle Optimizer Blog: Dynamic sampling and its impact on the Optimizer](#)
- Basic statistics collected during SQL compilation
 - Default level 2: if no statistics exist
 - Level 4: complex predicates, multiple predicates on same table

My Opinion

- Very effective for PeopleSoft
- Generally recommend level 4
 - This can cause parse problem in nVision with very large trees and 'use literal values' performance option.
- You may find some cases where explicit object optimizer statistics are still needed

RECOMMENDATION FOR STATISTICS ON TABLES USED FOR TEMPORARY WORKING STORAGE

- **Lock Statistics**

- to exclude them from maintenance windows statistics job
 - Temporary Table records
 - Normal Table records used for working storage (%AET, %TAO)

- **Delete Statistics**

- either the statistics are collected during the batch processes
- Or they are never collected and we can rely on dynamic statistics

Change DDL Model

- Use FORCE=>TRUE option in DBMS_STATS to override lock
- Consider using
 - GFCPSSTATS11 package
 - Or, Oracle's PSCBO_STATS package

Never

- Call DBMS_STATS directly in Application Engine because it implies a commit
 - use %UPDATESTATS instead.
 - The commit is why %UPDATESTATS is suppressed inside a DO WHILE loop

STATISTICS HISTORY

By Default

- Every time statistics are collected, old statistics are copied to statistics history tables (*WRI\$_OPTSTAT_%*)
- History is retained for 31 days

Problem

- Statistics on are collected some tables can be collected very frequently.
 - In some cases, thousands of times per process
 - *DBMS_STATS.GATHER_TABLE_STATS* is also responsible for purge.
 - I have seen it blocked by locks from other instances of itself.

Recommendation

- Disable Statistics History
- Set history retention to 0 days

```
exec dbms_stats.alter_stats_history_retention(0);
```

```
exec dbms_stats.purge_stats(SYSDATE);
```

Drawback

- Statistics History Retention is a global setting
 - Can't set it per table
- Lose the ability to determine table growth rates from the statistics history

'THE EVIL THAT MEN DO LIVES AFTER THEM; THE GOOD IS OFT INTERRED WITH THEIR BONES.'

- Huge numbers of SQL Statements processed.
- The optimizer does a perfectly good job with most of them
- Generally, performance problems occur when the optimizer makes poor choices on a few statements.

The optimizer has to be generic.

- It has to work with every application that has ever and will ever be written
- It is principally designed to work with applications that are written the way Oracle recommend.

However, PeopleSoft is not such an application.

- Not least because of its platform agnostic design.
- Including but no limited to:
 - Special Values
 - 0 and Space for NULL
 - Nearly all columns are Not Null
 - Unique Indexes instead of Primary Keys
 - Nullable columns in Unique Keys

HINTS

Hints are Directives to the Optimizer

- They mostly force certain operations
 - Eg. *FULL*, *INDEX*, *USE_NL*
- They can prevent certain transformations
 - Eg. *NO_EXPAND*, *NO_UNNEST*
- They can change optimizer environmental parameters for a statement
 - Eg. *OPT_PARAM('parallel_degree_policy','AUTO')*

Hint Paradox

- Why do we use hints?
 - To Improve Performance
- What happens to cost when we add a hint?
 - It usually increases*
- So the cost goes up but the runtime goes down!
- What does that say about optimizer cost?
 - There are problems with the cost model where the estimates are not correct, and it is often at the root of performance issues.

RECOMMENDATIONS

In those relatively few cases you need to intervene.

- Look for where the estimate cost diverges from the actual cost
- Make sure statistics are up to date
- Consider providing/removing histograms
- Consider indexes
- It may be necessary to use a hint to force or prevent optimizer behaviour.
- Hints are directives. That can be applied via
 - Outlines (deprecated in 11g)
 - SQL Baselines (match SQL_ID)
 - SQL Patches (match SQL_ID)
 - SQL Profiles (match FMS, but requires tuning pack)
 - force matching doesn't work well with a mixture of bind variables and literals
 - Directly in the code

Personal View

- I tend to favour putting hints into code in PeopleSoft
 - Note %SQLHint AE Macro
- I will not change COBOL source code
 - I have to use profiles on dynamic SQL (they never contain binds),
 - I will add hints into stored statements
- Sometimes, hints also have to be generated dynamically in dynamically generated SQL.

SOMETIMES, YOU NEED MORE CONTROL

Example: Time & Labor Calculation (TL_TIMEADMIN)

- Looped processing of batches of rules
- Statistics on a set of 20+ working storage tables collected for every loop
- Overhead of statistics collect can become dominant performance issue
- But dynamic statistics (optimizer dynamic sampling) not enough on some tables - need explicit statistics on some tables
- Don't want to change delivered code.

Change DDL Model

- Call custom PL/SQL Package from DDL Model
- Call DBMS_STATS from custom PL/SQL package
- Logic in package to determine whether to collect statistics

It is time to talk about

- *GFCPSSTATS11*
- *PSCBO_STATS*

GFCPSSTATS11

Written by Me

- PL/SQL Package
 - Superseded wrapper script designed for 10g
 - Rewritten and simplified for Oracle >=11g
- Metadata table (like a PeopleTools table)
 - Specifies statistics preferences
 - Also controls per PeopleTools record whether it
 - collect statistics
 - suppress statistics collection
 - delete statistics
 - or only collects statistics if stale
- Triggers to apply preferences
 - DDL Trigger on table creation
 - DML Trigger on metadata change
- Called from DDL Models 4 & 5
- Separately lock and delete statistics on temporary working tables

Advantages

- Leave Oracle maintenance window in place
- As far as is possible, manage statistics on PeopleSoft database in the same way as any other Oracle database in your estate.
- Control statistics collection behaviour declaratively with table statistics preferences
- Work with the database rather than fight against it.

Download

- Documentation: <http://www.go-faster.co.uk/docs.htm#Managing.Statistics.11g>
- From Github: <https://github.com/davidkurtz/gfcpsstats>

PSCBO_STATS

Written by Oracle

- Carlos Sierra, Oracle “Center of Excellence”
- PL/SQL Package
 - *Based on script for Siebel*
 - Dates back to Oracle 10
 - Enhanced for Oracle >=11g
- Control tables to specify statistics behaviour
- Delivered with PT8.57(?)
 - Documented in [PeopleBooks](#)
 - Download from Oracle Support Note [1322888.1](#)

Concept broadly similar to *GFCPSSTATS11*

- Both were developed completely independently
- To be fair, some of my original criticisms have been dealt with in the V2 release for Oracle 11g

Disadvantages

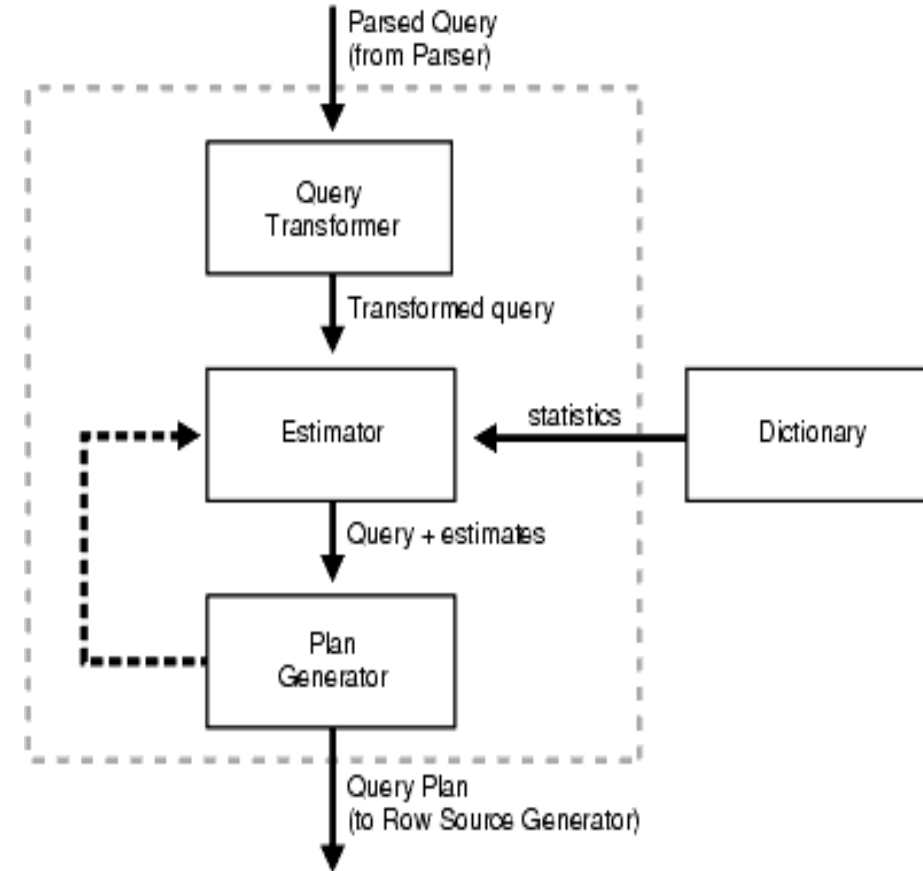
- Support *‘Note: Oracle does not certify or test the pscbo_stats utility for support purposes. Its use is considered to be a customization. If there is a problem or question with pscbo_stats, the Global Support Center will not be handling service requests regarding its use.’*
- No support for
 - Table Statistics Preferences
 - Extended Statistics
- Must always use package rather than *DBMS_STATS* directly.
- Replaces Default Oracle Maintenance Windows Stats Job

NOT ABOUT STATISTICS, BUT OPTIMIZER RELATED

ADAPTIVE OPTIMIZATION IN 12.1

Originally known as 'Cardinality Feedback'

- Additional statistics gathered
- Used to change decisions made by optimizer
- In a parse intensive system (such as PeopleSoft) this can become a significant overhead.



OPTIMIZER ADAPTIVE FEATURES SPLIT IN 12.2 (OR 12.1 WITH PATCHES)

OPTIMIZER_ADAPTIVE_PLANS=TRUE

- Very Beneficial
 - Nested loop/Hash join selection
- Less Important in PeopleSoft
 - Adaptive Parallel distribution
 - Star Transformation Bitmap Pruning

OPTIMIZER_ADAPTIVE_STATISTICS=FALSE

- SQL Plan Directives
- Statistic Feedback for Joins
- Performance Feedback
- Adaptive Sampling for Parallel Execution
- Enabling this causes performance problems at compile time
- Effectively TRUE in 12.1 when Adaptive Features enabled

(see [PeopleSoft and Adaptive Query Optimization in Oracle 12c](#))

(see [Oracle Optimizer Blog: Optimizer Adaptive Features in Oracle Database 12c Release 2](#))

ADAPTIVE OPTIMIZATION RECOMMENDATIONS

In 12.1 (without patches)

- Proactive Bundle for Patches **22652097** and **21171382** and follow 12.2 advice
- See <https://mikedietchde.com/2017/07/06/adpative-features-patches-oracle-peoplesoft/>
- Otherwise
`OPTIMIZER_ADAPTIVE_FEATURES=FALSE`
- Set `OPTIMIZER_DYNAMIC_SAMPLING` to a non-default value
 - I recommend 4
 - Or 1 if parse problems in nVision

In 12.2 (or 12.1 with patches)

- Use the Defaults
 - `OPTIMIZER_ADAPTIVE_STATISTICS=FALSE`
 - `OPTIMIZER_ADAPTIVE_PLANS=TRUE`
- Set `OPTIMIZER_DYNAMIC_SAMPLING=4`
 - Or leave default if parse problems in nVision

SQL PLAN DIRECTIVES

SQL Plan Directives

- Still generated but not used if Optimizer Adaptive Statistics disabled
- Matched on SQL_ID, so not effective with
 - dynamically generated literal values
 - Other dynamically generated code
- Will continue to build up over time
- We experienced a progressive increase in parse CPU in 12.1.0.2.
- Resolved by dropping all directives, but couldn't reproduce on another environment!

Recommendation

- Disable directive generation by setting
 - ***_sql_plan_directive_mgmt_control=0***
 - (or ***_optimizer_dmdir_usage_control=0***)
- Drop Existing Plan Directives
 - Expect 3s / directive

```
set serveroutput on timi on
DECLARE
  l_count INTEGER := 0;
BEGIN
  FOR i IN (select * from dba_sql_plan_directives) LOOP
    dbms_output.put_line('Dropping SPD '||i.directive_id);
    dbms_spd.drop_sql_plan_directive(i.directive_id);
    l_count := l_count+1;
  END LOOP;
  dbms_output.put_line(l_count||' directives dropped');
END;
/
```

DATABASE INITIALISATION PARAMETERS

_unnest_subquery=FALSE

- Required at Install since Oracle 9i
- **It is critically important that this parameter is set**
 - But I still find PeopleSoft environments where this hasn't been set
- Prevents correlated sub-queries being unnested into in-line views.
 - The optimizer does this because it miscosts the cardinality of correlated columns.
 - If not set, especially severe performance effect in HR, but all products affected

_gby_hash_aggregation_enabled=FALSE

- Recommended at Install by PeopleSoft/Oracle since Oracle 9i
 - But only because some SQL statements do not have ORDER BY clause.
 - This setting prevents GROUP BY and DISTINCT being done with a HASH if cheaper and reverts to the original sort method.
 - When using a sort to perform a GROUP BY or DISTINCT on a non-partitioned object, the rows are usually returned in that order.
 - But a different result set order can occur if the plan changes even if this parameter is set
- **I do not think this parameter is needed** because modern PeopleSoft applications have been updated to use ORDER BY
- SQL92 standard says that row order is arbitrary without an ORDER BY clause.
- If you care about order, add an ORDER BY clause.

CONCLUSION

CONCLUSION

Do it the Oracle way, work with the database as far as is possible.

- Collect optimizer statistics in the maintenance window
- Use GFCPSSTATS11 instead of PSCBO_STATS
 - To manage table statistic preferences
 - Finer grain control of when to collect statistics during processing
 - Data driven
- Disable statistics history retention
 - (unfortunate but necessary)
- Use Adaptive Plans but not Adaptive Statistics from 12c
 - 12.1 patches to backport 12.2 features
 - 12.2 default parameters
 - Disable and purge SQL Plan Directives
 - Set the 'required at install' parameters
 - (except *_gby_hash_aggregation_enabled*)