

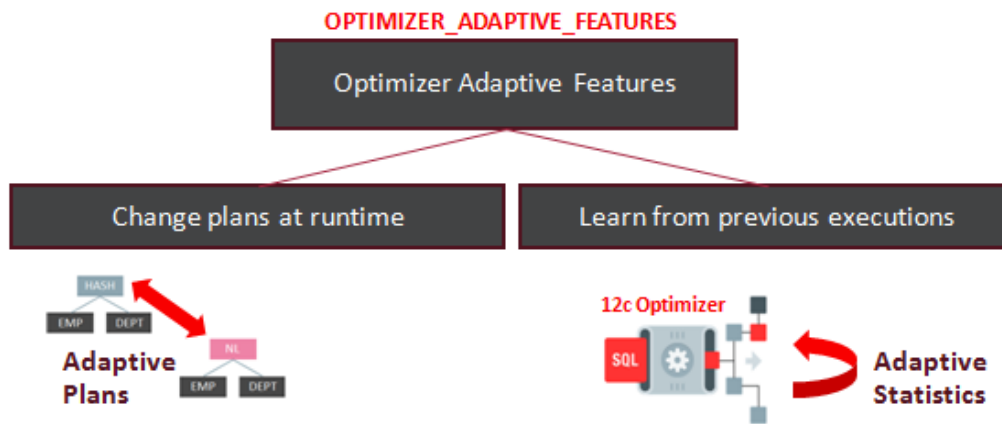
Optimizer Adaptive Features in the Exadata Express Cloud Service

Introduction

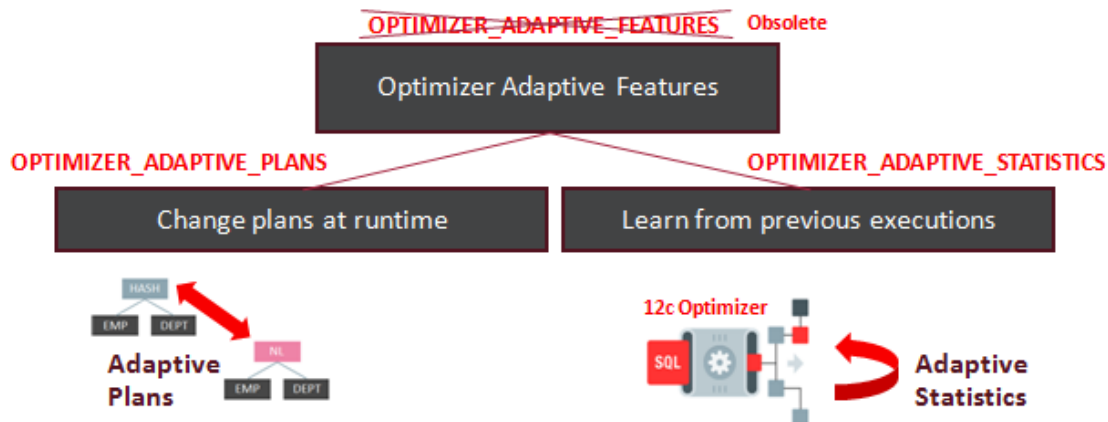
In the Exadata Express Cloud Service (Oracle Database 12c Release 2) we have changed the way optimizer adaptive features can be controlled. In this post, I'll present what's changed and give you guidance on how to you can choose what settings are best for your systems.

What's Changed

In Oracle Database 12c Release 1, the database parameter `OPTIMIZER_ADAPTIVE_FEATURES` controls all of the adaptive features like this:



In Oracle Database 12c Release 2, this parameter has been made obsolete and replaced with two new parameters that control adaptive plans and adaptive statistics separately, like this:



What are the Optimizer Adaptive Features?

Adaptive plans and adaptive statistics are umbrella terms for a number of optimizer features. The following tables summarize them.

These features are enabled by default:

OPTIMIZER_ADAPTIVE_PLANS – Default TRUE	Description
Nested loop join/Hash join selection	The optimizer chooses between nested loops or hash joins at query runtime.
Adaptive parallel distribution method	The parallel distribution method is determined at runtime.
Star transformation bitmap pruning	Certain bitmap indexes may be removed from the SQL execution plan at runtime if selectivity is significantly poorer than the estimate.

These features are disabled by default:

OPTIMIZER_ADAPTIVE_STATISTICS – Default FALSE	Description
SQL plan directives	SQL plan directives are created and used to adapt SQL execution plans.
Statistics feedback for joins	Cardinality from table joins is used to improve SQL execution plans.
Performance feedback	Improves the degree of parallelism chosen when PARALLEL_DEGREE_POLICY is set to ADAPTIVE
Adaptive dynamic sampling for parallel execution	Dynamic statistics are gathered for certain parallel queries to improve cardinality estimates. The sample size is determined automatically.

SQL Plan Directives

Setting OPTIMIZER_ADAPTIVE_STATISTICS to FALSE does not disable the *creation* of SQL plan directives. Instead, this setting prevents them being used with dynamic statistics to influence SQL execution plans.

Statistics Feedback

Setting OPTIMIZER_ADAPTIVE_STATISTICS to FALSE disables Oracle Database 12c statistics feedback for *joins*. Note that statistics feedback for cardinality (introduced in Oracle Database 11g as *cardinality feedback*) is always enabled.

Adaptive Dynamic Sampling

If OPTIMIZER_ADAPTIVE_STATISTICS is set to TRUE then dynamic statistics using an adaptive sample size will kick in for certain parallel queries if [OPTIMIZER_DYNAMIC_SAMPLING](#) is 2 (the default value). If

OPTIMIZER_ADAPTIVE_STATISTICS is set to FALSE, then adaptive dynamic sampling will not be used for parallel execution.

Changes to Auto Creation of Column Group Statistics

In Oracle Database 12c Release 1, SQL Plan Directives trigger the creation of column group statistics when statistics are gathered. You can see the extended statistics created automatically using a query like this:

```
select owner,
       table_name,
       extension,
       extension_name
from dba_stat_extensions
where creator = 'SYSTEM'
order by owner,table_name,extension_name;
```

We received feedback that some DBAs wanted to be able to control this feature so in Oracle Database 12c Release 2, automatic column group statistics creation is controlled by a DBMS_STATS preference AUTO_STAT_EXTENSIONS. By default, the preference is **OFF** so that extended statistics are **not** created automatically. You can re-enable this feature using:

```
EXEC DBMS_STATS.SET_GLOBAL_PREFS('AUTO_STAT_EXTENSIONS', 'ON')
```

Choosing What Settings to Use

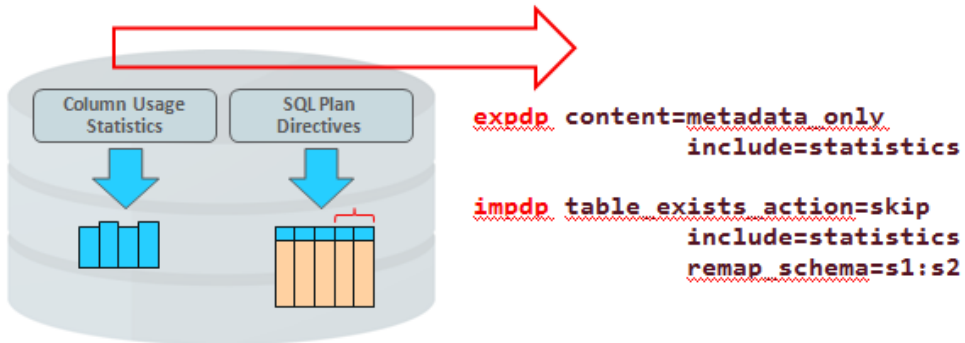
The defaults have been chosen on the basis of what works best for widest range of workloads, so it is recommended that you start with these for most systems. If you're already using all the adaptive features in Oracle Database 12c Release 1, and you want the same in the Exadata Express Cloud Service, then override the defaults by setting the parameter OPTIMIZER_ADAPTIVE_STATISTICS to TRUE and the preference AUTO_STAT_EXTENSIONS preference to ON.

We have received very positive feedback from customer with respect to adaptive plans. It has proved to be very able at improving system performance for a huge range of workloads. In other words, adapting the execution plan at runtime is likely to benefit all types of systems so we elected to set the default to TRUE. Of course, if you prefer not to use it initially then you can set the parameter to FALSE.

Adaptive statistics features are more useful in environments where queries are longer-running and where data distributions and schemas are very complex. In systems like this, where query execution times are longer, it makes especially good sense for the optimizer to invest time and resources to improve the overall quality of the execution plans. The default value of this parameter is FALSE and can be set to TRUE in these environments.

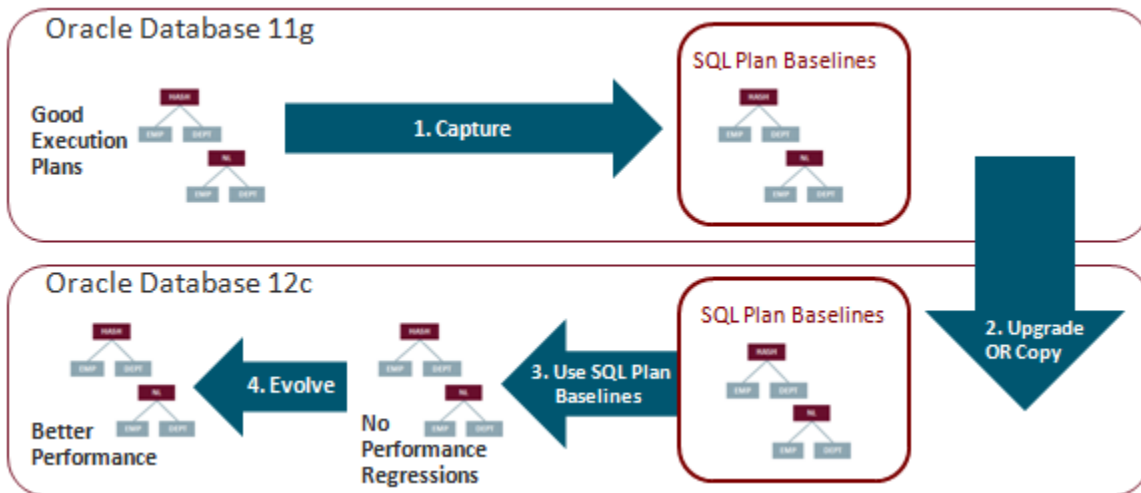
Copying Optimizer Metadata

Use data pump if you need to copy optimizer metadata from one database to another.

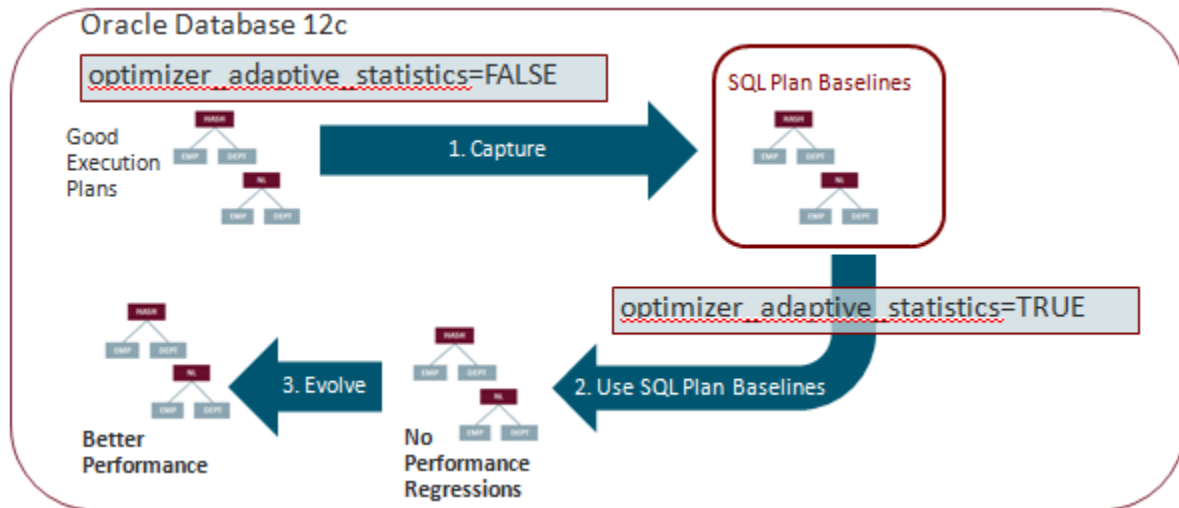


Guaranteed Plan Stability

Use SQL plan management when upgrading to guarantee plan stability:



Use SQL plan management to mitigate the risk of change:



Oracle Database 12c Release 1

If you want the new adaptive parameters in Oracle Database 12c Release 1 request patch for bug# 22652097.

To control auto column group creation using DBMS_STATS preference AUTO_STAT_EXTENSIONS, apply patch for bug# 21171382.

Further Information

Oracle has made improvements to the way adaptive features are controlled. If you want more information, particularly if you are currently using or upgrading to Oracle Database 12c Release 1, then check out the blogs and MOS:

- Oracle Optimizer Blog: <http://blogs.oracle.com/optimizer>
- Upgrade Blog: <http://blogs.oracle.com/UPGRADE>
- Recommendations for Adaptive Features in Oracle Database 12c Release 1 (12.1) (Doc ID 2187449.1)

If you need to find out more about the adaptive features, there is a [white paper](#) for Oracle Database 12c Release 1. An update to this paper is coming up.