



Product Documentation

Embarcadero® DB Optimizer™

New Features Guide

Version 3.8/XE5

Published February, 2014

© 2014 Embarcadero Technologies, Inc. Embarcadero, the Embarcadero Technologies logos, and all other Embarcadero Technologies product or service names are trademarks or registered trademarks of Embarcadero Technologies, Inc. All other trademarks are property of their respective owners.

This software/documentation contains proprietary information of Embarcadero Technologies, Inc.; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

Embarcadero Technologies, Inc. is a leading provider of award-winning tools for application developers and database professionals so they can design systems right, build them faster and run them better, regardless of their platform or programming language. Ninety of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero products to increase productivity, reduce costs, simplify change management and compliance, and accelerate innovation. Founded in 1993, Embarcadero is headquartered in San Francisco, with offices located around the world. To learn more, please visit <http://www.embarcadero.com>.

April 2, 2014

Contents

New Features	4
Platform/version Support Updates	4
Tuning Updates	4

NEW FEATURES

New features for this release fall into the following categories:

- [Platform/version Support Updates](#)
- [Tuning Updates](#).

PLATFORM/VERSION SUPPORT UPDATES

This release introduces execution, tuning, and profiling support for the following DBMS versions:

- Oracle 12c
- SQL Server 2012
- Sybase ASE 15.7.

TUNING UPDATES

The following topics describe SQL Tuner updates for this release.

SQL Server Hints

The following table lists the SQL Server hints, newly available for this release:

Hint	Description
IGNORE_NONCLUSTERED_COLUMNSTORE_INDEX	Lets you disable use of a nonclustered xVelocity memory optimized columnstore index.
SPATIAL_WINDOW_MAX_CELLS	Specifies the maximum number (1 - 8192) of cells to use when tessellating a geometry or geography object.

DB Optimizer no longer provides support for the **FASTFIRSTROW** hint.

Oracle 12c Hints

The following table lists the Oracle 12c hints, newly available for this release:

Hint	Description
GATHER_OPTIMIZER_STATISTICS	Instructs the optimizer to enable statistics gathering during CREATE TABLE... AS SELECT and INSERT INTO ... SELECT bulk loads.
NO_GATHER_OPTIMIZER_STATISTICS	Instructs the optimizer to disable statistics gathering during CREATE TABLE... AS SELECT and INSERT INTO ... SELECT bulk loads.
NO_PQ_CONCURRENT_UNION	Instructs the optimizer to disable concurrent processing of UNION and UNION ALL operations.

Hint	Description
NO_PQ_SKEW	Advises the optimizer that the distribution of values of the join keys for a parallel join is not skewed.
NO_USE_CUBE	Instructs the optimizer to exclude cube joins when joining each specified table to another row source using the specified table as the inner table.
PQ_CONCURRENT_UNION	Instructs the optimizer to enable concurrent processing of UNION and UNION ALL operations.
PQ_FILTER	This hint takes one of four arguments: SERIAL: Process rows serially on the left and right sides of the filter. Use this option when the overhead of parallelization is too costly for the query. NONE: Process rows in parallel on the left and right sides of the filter. Use this option when there is no skew in the distribution of the data on the left side of the filter and you want to avoid distribution of the left side. HASH: Process rows in parallel on the left side of the filter using a hash distribution and serially on the right side of the filter. Use this option when there is no skew in the distribution of data on the left side of the filter. RANDOM: Process rows in parallel on the left side of the filter using a random distribution and serially on the right side of the filter. Use this option when there is skew in the distribution of data on the left side of the filter.
PQ_SKEW	Advises the optimizer that the distribution of values of the join keys for a parallel join is highly skewed.
USE_CUBE	When the right-hand side of the join is a cube, this hint instructs the optimizer to join each specified table with another row source using a cube join. If the optimizer decides not to use the cube join based on statistical analysis, you can use USE_CUBE to override that decision.

Oracle 11g Hints

The following table lists the Oracle 11g hints, newly available for this release:

Hint	Description
APPEND_VALUES	Instructs the optimizer to use direct-path INSERT with the VALUES clause.
CHANGE_DUPKEY_ERROR_INDEX	Unambiguously identifies a unique key violation for a specified index or set of columns.
RETRY_ON_ROW_CHANGE	Retries the operation when the ORA_ROWSCN for one or more rows in a set has changed from the time the set of rows to be modified is determined, to the time the block is actually modified.