

Database In-Memory

Enabling Real-Time Analytics

Abstract

Oracle Database In-Memory introduces a new in-memory columnar format that is compatible with all existing applications. Join Oracle Product Management to learn in detail what motivated Oracle to develop this new technology and how it works alongside the traditional row format to provide a unique “dual format” architecture. This session will explore how Database In-Memory enables "Real-Time Analytics" by allowing analytic style queries to be run directly on source data. It will also explore the key performance enhancing features of Database In-Memory including in-memory scans, predicate push down, hash joins with bloom filters, vector group by aggregation, and new high-performance features like In-Memory Expressions and Join Groups.

It will then examine what aspects of a SQL query benefit from Database In-Memory and compare some actual SQL statement executions to demonstrate where Database In-Memory provides benefit and explain how Database In-Memory provides that benefit. We will then conclude by describing what types of workload benefit from Database In-Memory and what customers are doing with Database In-Memory.

Agenda

- What is Database In-Memory?
- What workload benefits from Database In-Memory
- Getting started
- How does it work?
- Population
- Why is it faster?
- What about OLTP?
- How does it work with other database features?
- How customers are using Database In-Memory

Speaker



Andy Rivenes is a Product Manager at Oracle for Database In-Memory. Prior to that Andy was a Product Manager for Heat Map, Automatic Data Optimization, SecureFiles and DBFS. Andy has spent a large part of his career managing Oracle databases and he now spends his time helping customers implement Database In-Memory as well as providing information about Database In-Memory through blogs, conferences, and customer visits. You can visit the Database In-Memory blog at <https://blogs.oracle.com/in-memory> or follow his tweets as @TheInMemoryGuy.