

# Oracle Database In-Memory

## *By Example*



**Andy Rivenes**  
Database In-Memory Product Management  
**Oracle** Corporation

Email: [andy.rivenes@oracle.com](mailto:andy.rivenes@oracle.com)  
Twitter: @TheInMemoryGuy  
Blog: [blogs.oracle.com/in-memory](https://blogs.oracle.com/in-memory)



# Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

# What Is Database In-Memory



# Row Format Databases vs. Column Format Databases

Rows Stored  
Contiguously



- **Transactions** run faster on row format
  - Example: Query or Insert a sales order
  - Fast processing few rows, many columns

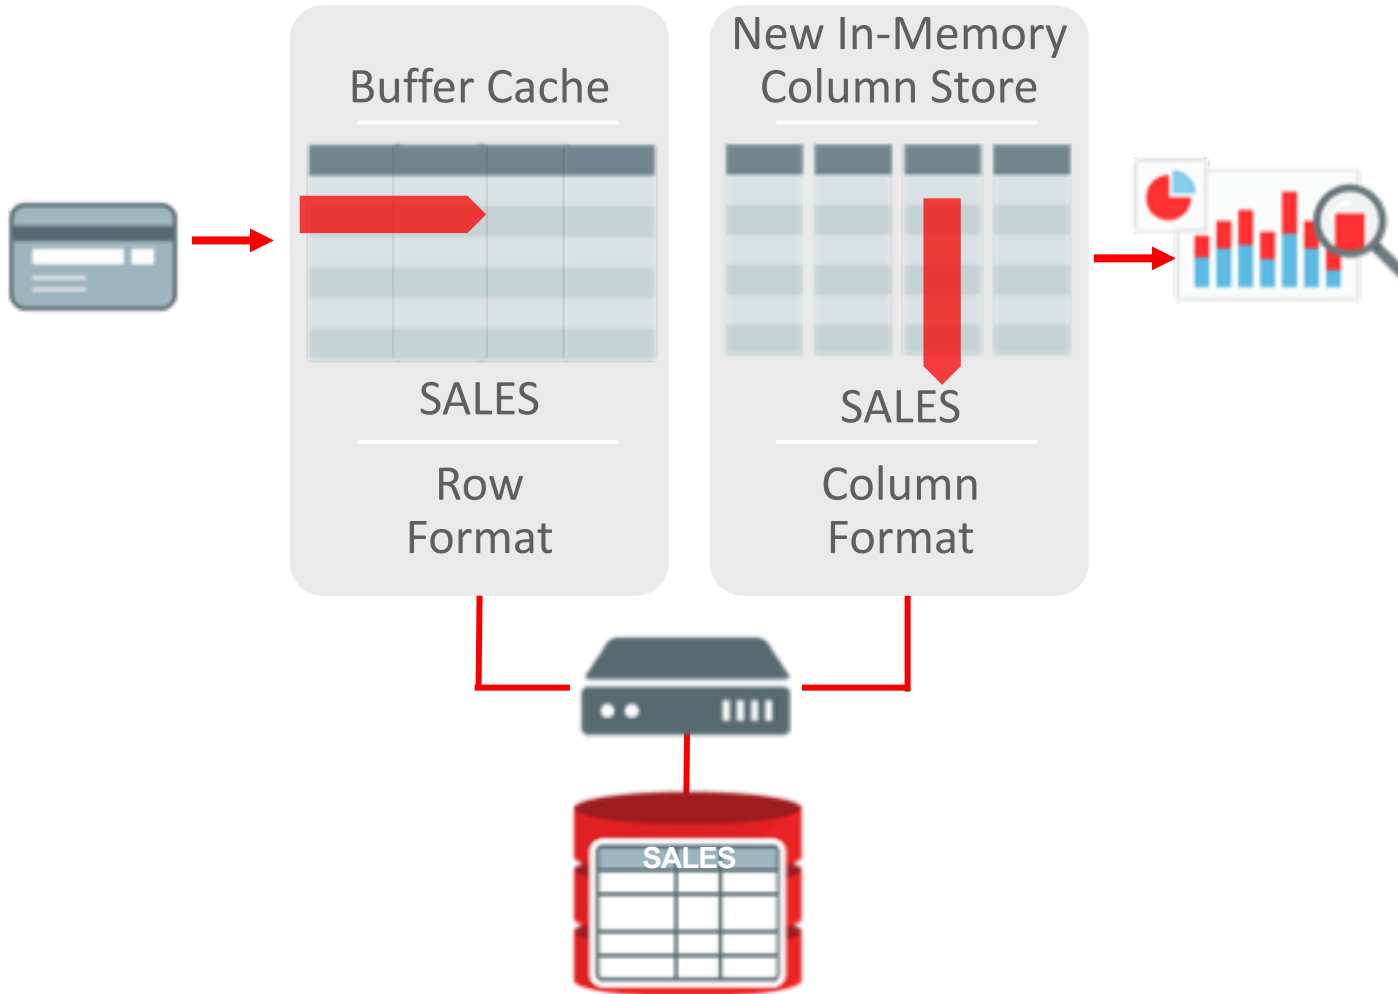
Columns  
Stored  
Contiguously



- **Analytics** run faster on column format
  - Example : Report on sales totals by region
  - Fast accessing few columns, many rows

Until Now Must Choose One Format and Suffer Tradeoffs

# Breakthrough: Dual Format Database



- **BOTH** row and column formats for same table
- Simultaneously active and transactionally consistent
- Analytics & reporting use new in-memory Column format
- OLTP uses proven row format

# Oracle In-Memory: Simple to Implement

## 1. Configure Memory Capacity

- `inmemory_size = XXX GB`

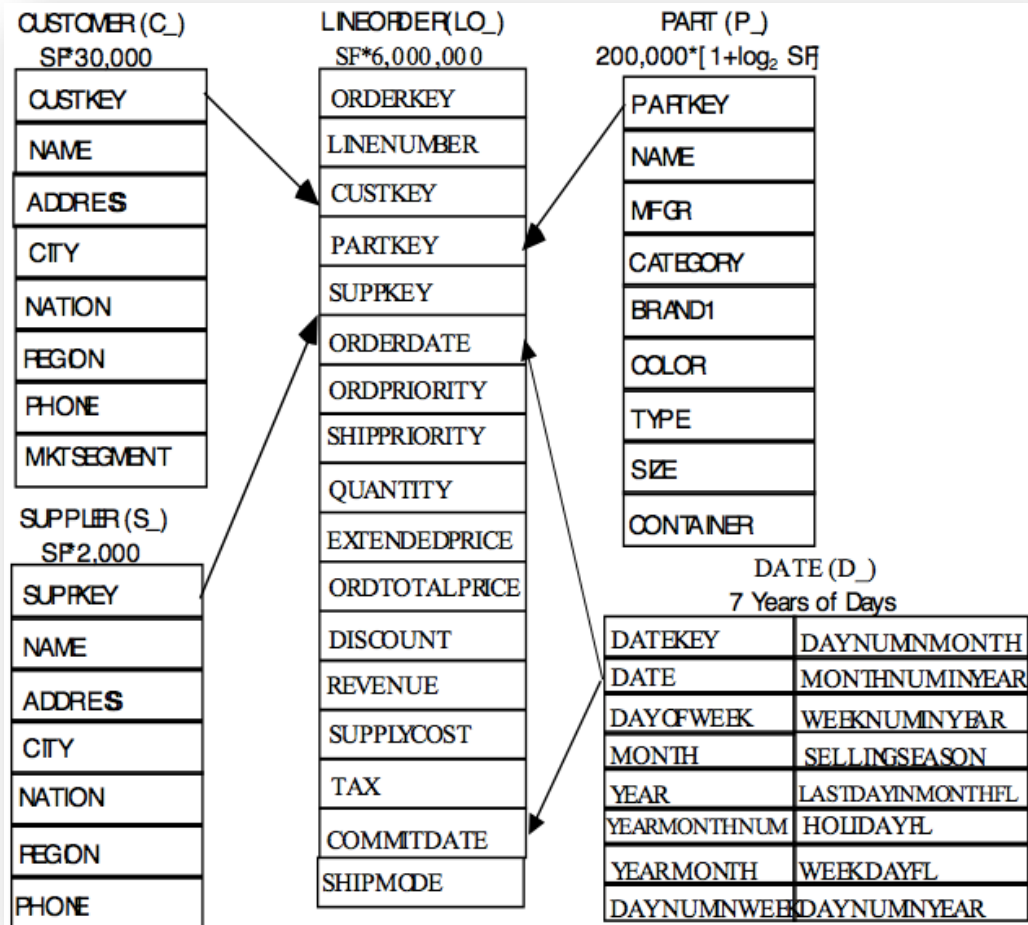
## 2. Configure tables or partitions to be in memory

- `alter table | partition ... inmemory;`

## 3. Later drop analytic indexes to speed up OLTP

# Database In-Memory By Example Details

# In-Memory By Example - Schema Overview



- Star schema based on TPCH benchmark
- Lineitem & Orders tables combined to create fact table
- Schema built with a 3GB scale factor



# In-Memory By Example

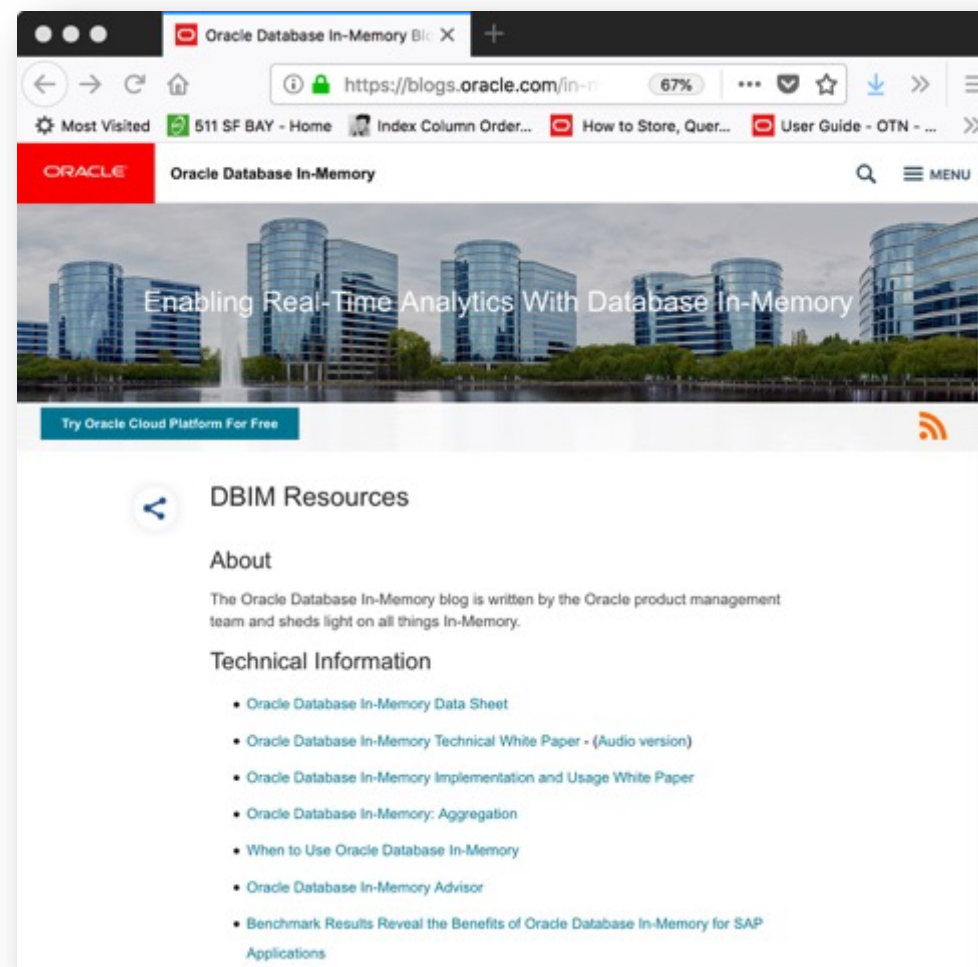
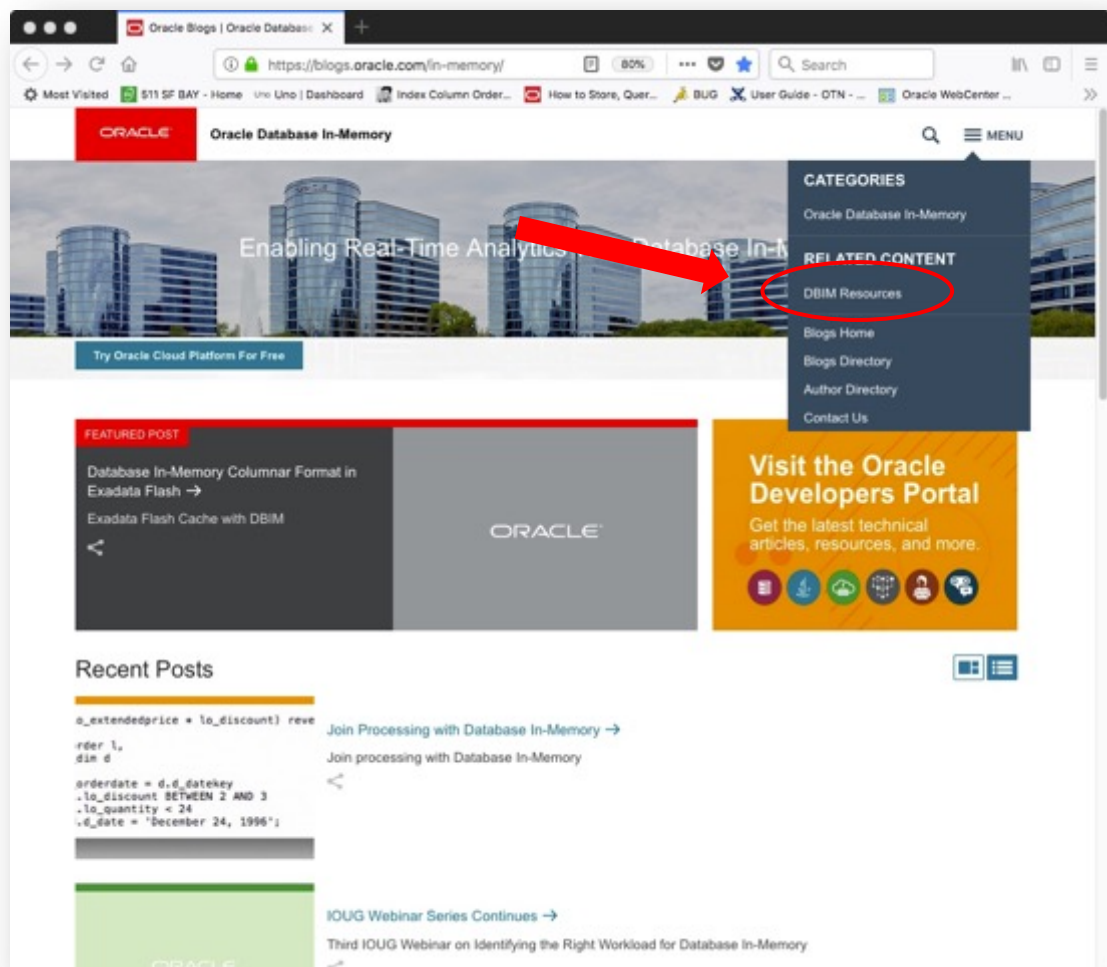
## Demonstration Details

- Virtual Box on my laptop – 12GB memory size
- Set of SQL scripts to show how Database In-Memory works
- SGA size = 10000M
- IM column store size = 5504M
- Keep pool size = 4000M (cache row store tables)
- Tables have "cache" attribute set
- Each script connects as a new session

# Where Can You Get More Information



# https://blogs.oracle.com/in-memory/dbim-resources



# Additional Resources



## Database In-Memory Information

[Database In-Memory Blog](#)

[oracle.com – Database In-Memory](#)

[Database In-Memory YouTube Channel](#)

[Ask TOM Database In-Memory Office Hours](#)

[Database In-Memory Guide \(Documentation\)](#)

## Join the Conversation

 [https://twitter.com/db\\_inmemory](https://twitter.com/db_inmemory)

 <https://twitter.com/TheInMemoryGuy>

 <http://www.oracle.com/goto/dbim.html>