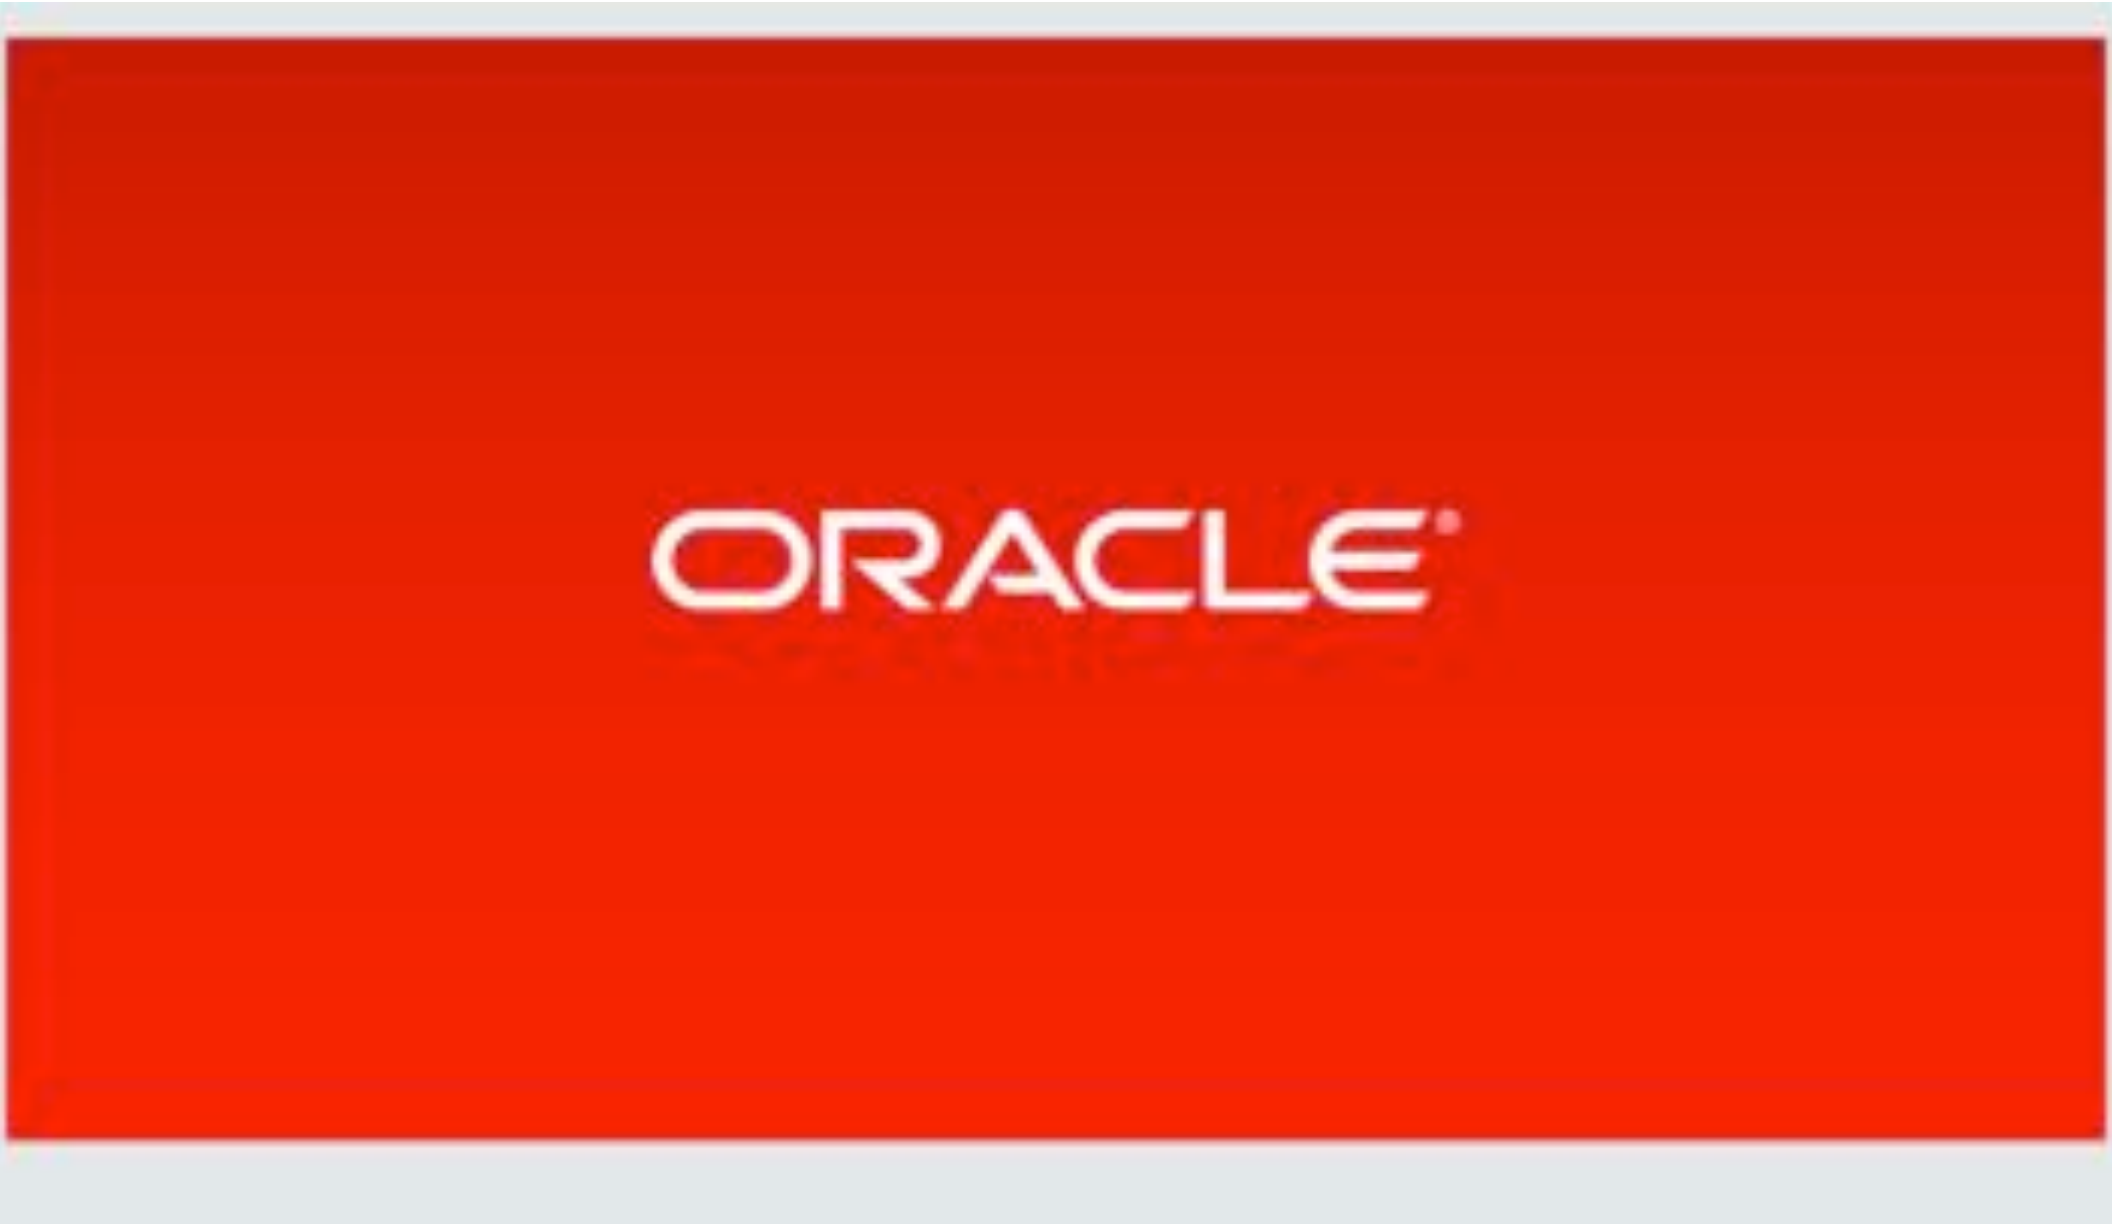
ORACLE

Comparing and Contrasting Oracles In-Memory Technologies

Gerald Venzl

Principal Sales Consultant

Oracle UK

ORACLE

Program Agenda

1. Overview
2. Oracle TimesTen
3. Oracle Database In-Memory Option
4. Customers

Program Agenda

1. Overview
2. Oracle TimesTen
3. Oracle Database In-Memory Option
4. Customers

Overview

Overview

Overview

Program Agenda

1. Overview
2. Oracle TimesTen
3. Oracle Database In-Memory Option
4. Customers

Oracle TimesTen History

- Founded in HP Labs in 1995
- Spun off into a separate startup in 1996
- Released in 1998 as the first commercial In-Memory RDBMS
- Acquired by Oracle in 2005

ORACLE

Oracle TimesTen In-Memory Database

Microseconds Response Time in Application Tier

- Relational Database
  - Pure in-memory
  - ACID compliant
  - Standard SQL
  - Entire database in DRAM

- Persistent and Recoverable
  - Database and Transaction logs persisted on disk and flash storage (local to TimesTen)

- Extremely Fast
  - Shared library approach
  - Microseconds response time
  - Very high throughput

- Compatible with Oracle Database
  - Data types, PL/SQL, OCI, ODP.NET, PHP, R
  - Integrated with RAC, Data Guard, Enterprise Manager, SQL Developer, etc.

Real-Time Transactional Replication

High Availability and Disaster Recovery

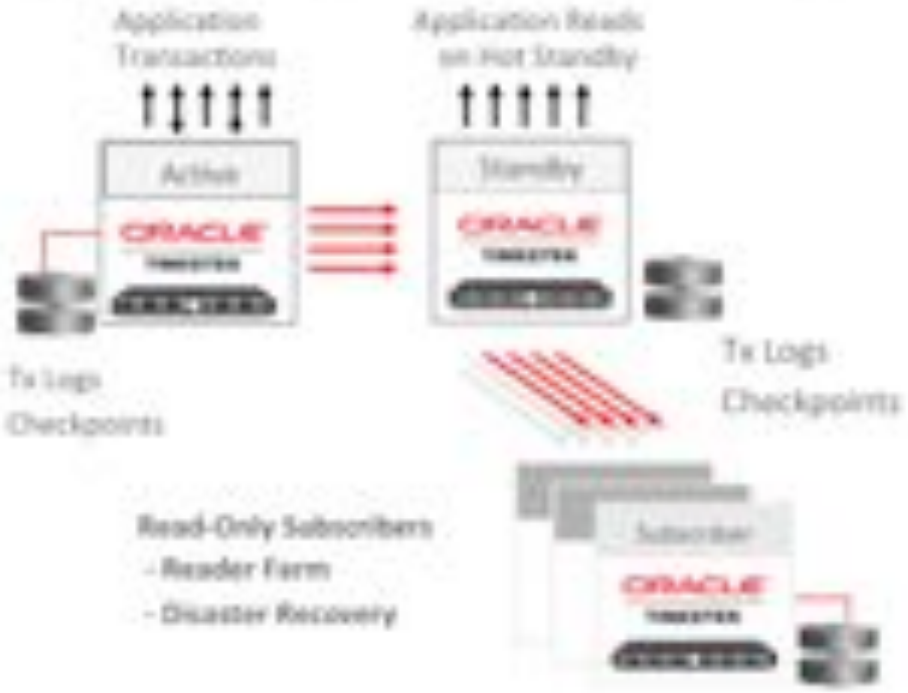

- High performance
  - Synchronous / Asynchronous
  - Parallel send of log streams
  - Parallel apply of changes on Standby and Subscribers
- HA and DR support
- Online upgrades
  - No downtime
  - Cross-version replication
- Integration with Oracle Clusterware

Oracle TimesTen Deployment Mode

Standalone

Application-Tier Cache

Adaptive In-Memory Cache for Exalytics

ORACLE

Program Agenda

- Overview
- Oracle TimesTen
- **Oracle Database In-Memory Option**
- Customers

Oracle Database In-Memory Option

- Released with latest Oracle Database version 12.1.0.2
- Cost option on top of Oracle Database Enterprise Edition
- Holds relational data entirely in memory
- Represents relational data in a columnar format

Oracle Database In-Memory Option Goals

- **100x** Faster Queries: Real-Time Analytics
  - Instantaneous Queries on OLTP Database or Data Warehouse
- Transparent: No application changes
  - Simple to Implement

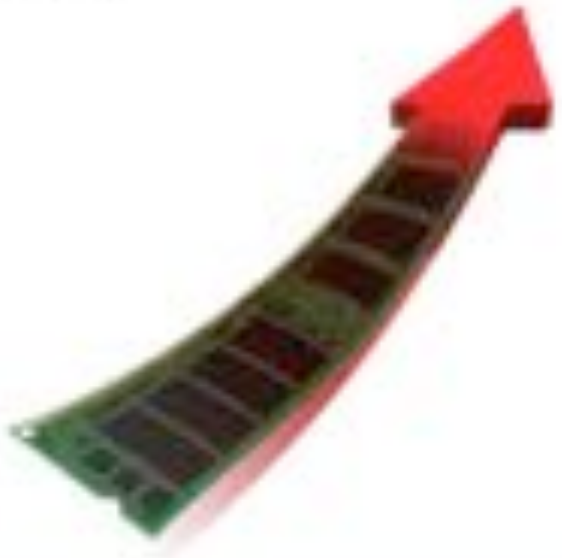

ORACLE

Optimizing Transaction and Query Performance

Row Format Databases versus Column Format Databases

Row

- Transactions run faster on row format
  - Insert or query a sales order
  - Fast processing few rows, many columns

Column

- 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

ORACLE

Breakthrough: Dual Format In-Memory Database

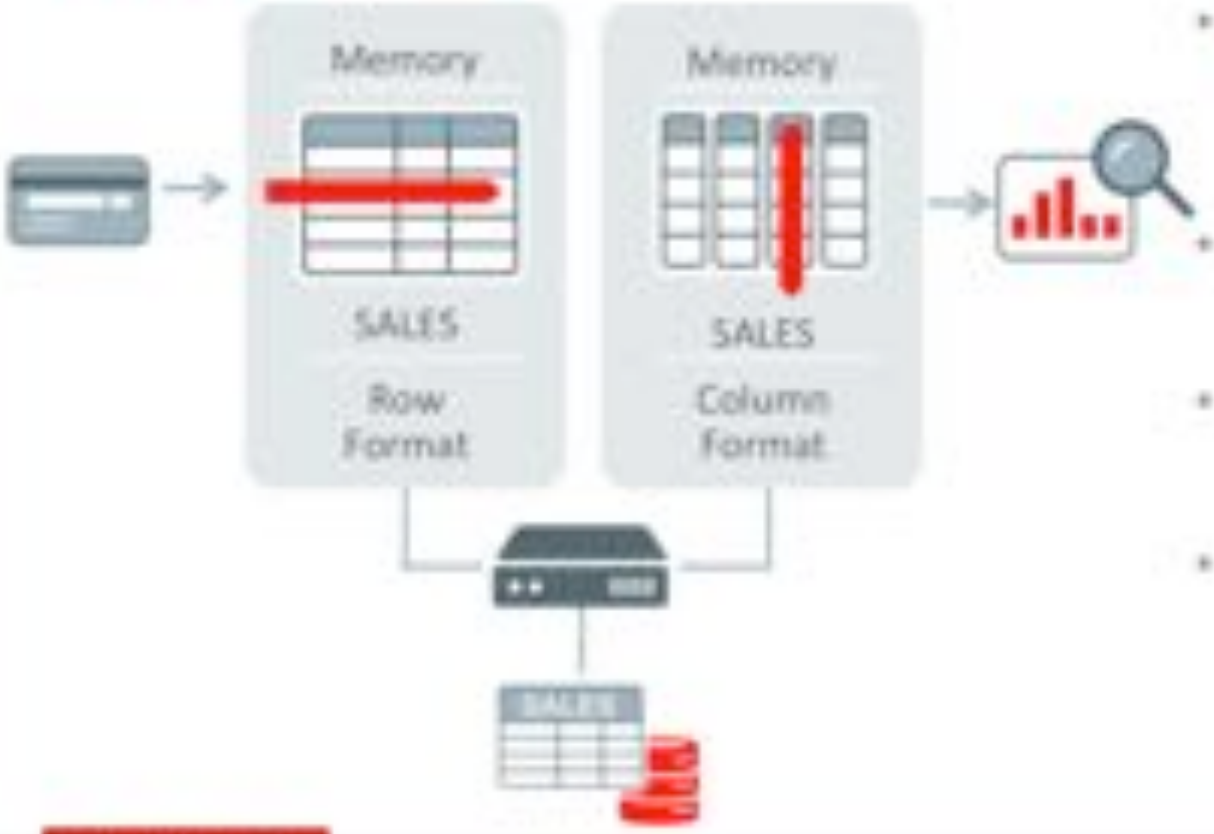

- **BOTH** row and column in-memory formats for same table
- Simultaneously active and transactionally consistent
- Analytics & reporting use new in-memory Column format
- OLTP uses proven row format

Why is an In-Memory scan faster than the buffer cache?

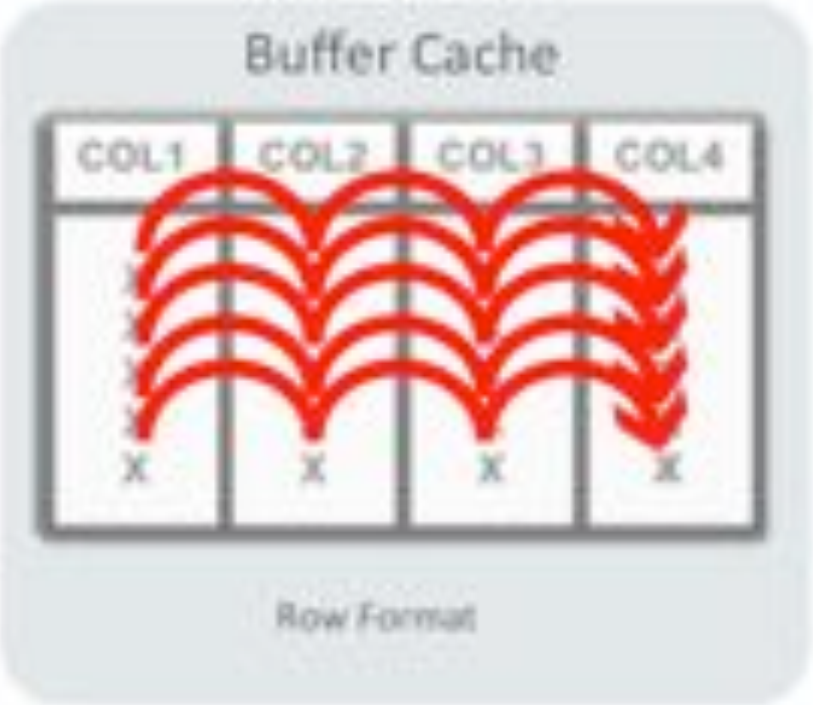


```
SELECT COL4 FROM MYTABLE;
```

Why is an In-Memory scan faster than the buffer cache?

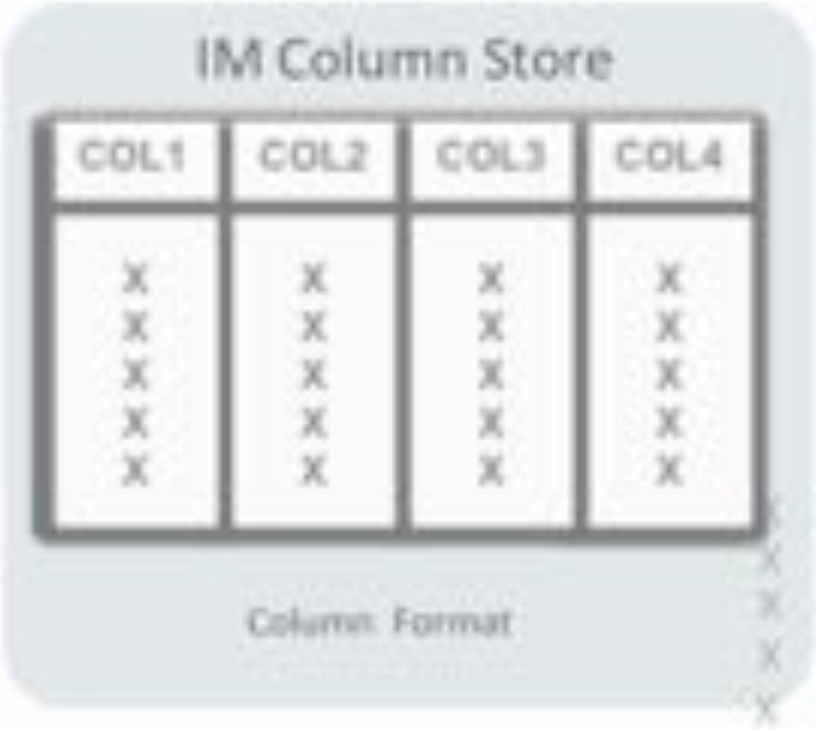


```
SELECT COL4 FROM MYTABLE;
```

Program Agenda

1. Overview
2. Oracle TimesTen
3. Oracle Database In-Memory Option
4. Customers

Oracle TimesTen Customers

Top 5 reasons customers select TimesTen

- Microsecond response time
- Consistent response time
- Handles workload peaks
- Provides uninterrupted service
- Requires minimal application changes

ORACLE

Oracle In-Memory Option Customers

Including beta testing customers

Top 5 reasons customers select Oracle DB In-Memory Option

- Speed up analytical queries
- Real-Time analytics on transactional DB
- Real-Time reporting
- Speed up ETL processing
- Storage saving by dropping indexes

ORACLE

Capabilities Matrix

	Objects	Relational Data	Parallel processing	Ultra Low latency (no network hops)	Event driven	Reporting driven	Txn	Query Lang	Joins
Coherence	●	○	●	◐	●	○	◐	◐	○
TimesTen	○	●	○	●	●	○	●	●	●
In-Memory Option	○	●	●	○	◐	●	●	●	●

Capabilities Matrix

	Coherence	TimesTen	DB In-Mem
Objects			
Relational Data			
Parallel processing			
Ultra Low latency			
Event driven			
Reporting driven			
Transactions			
Query language			
Joins			

Hardware and Software
Engineered to Work Together

ORACLE

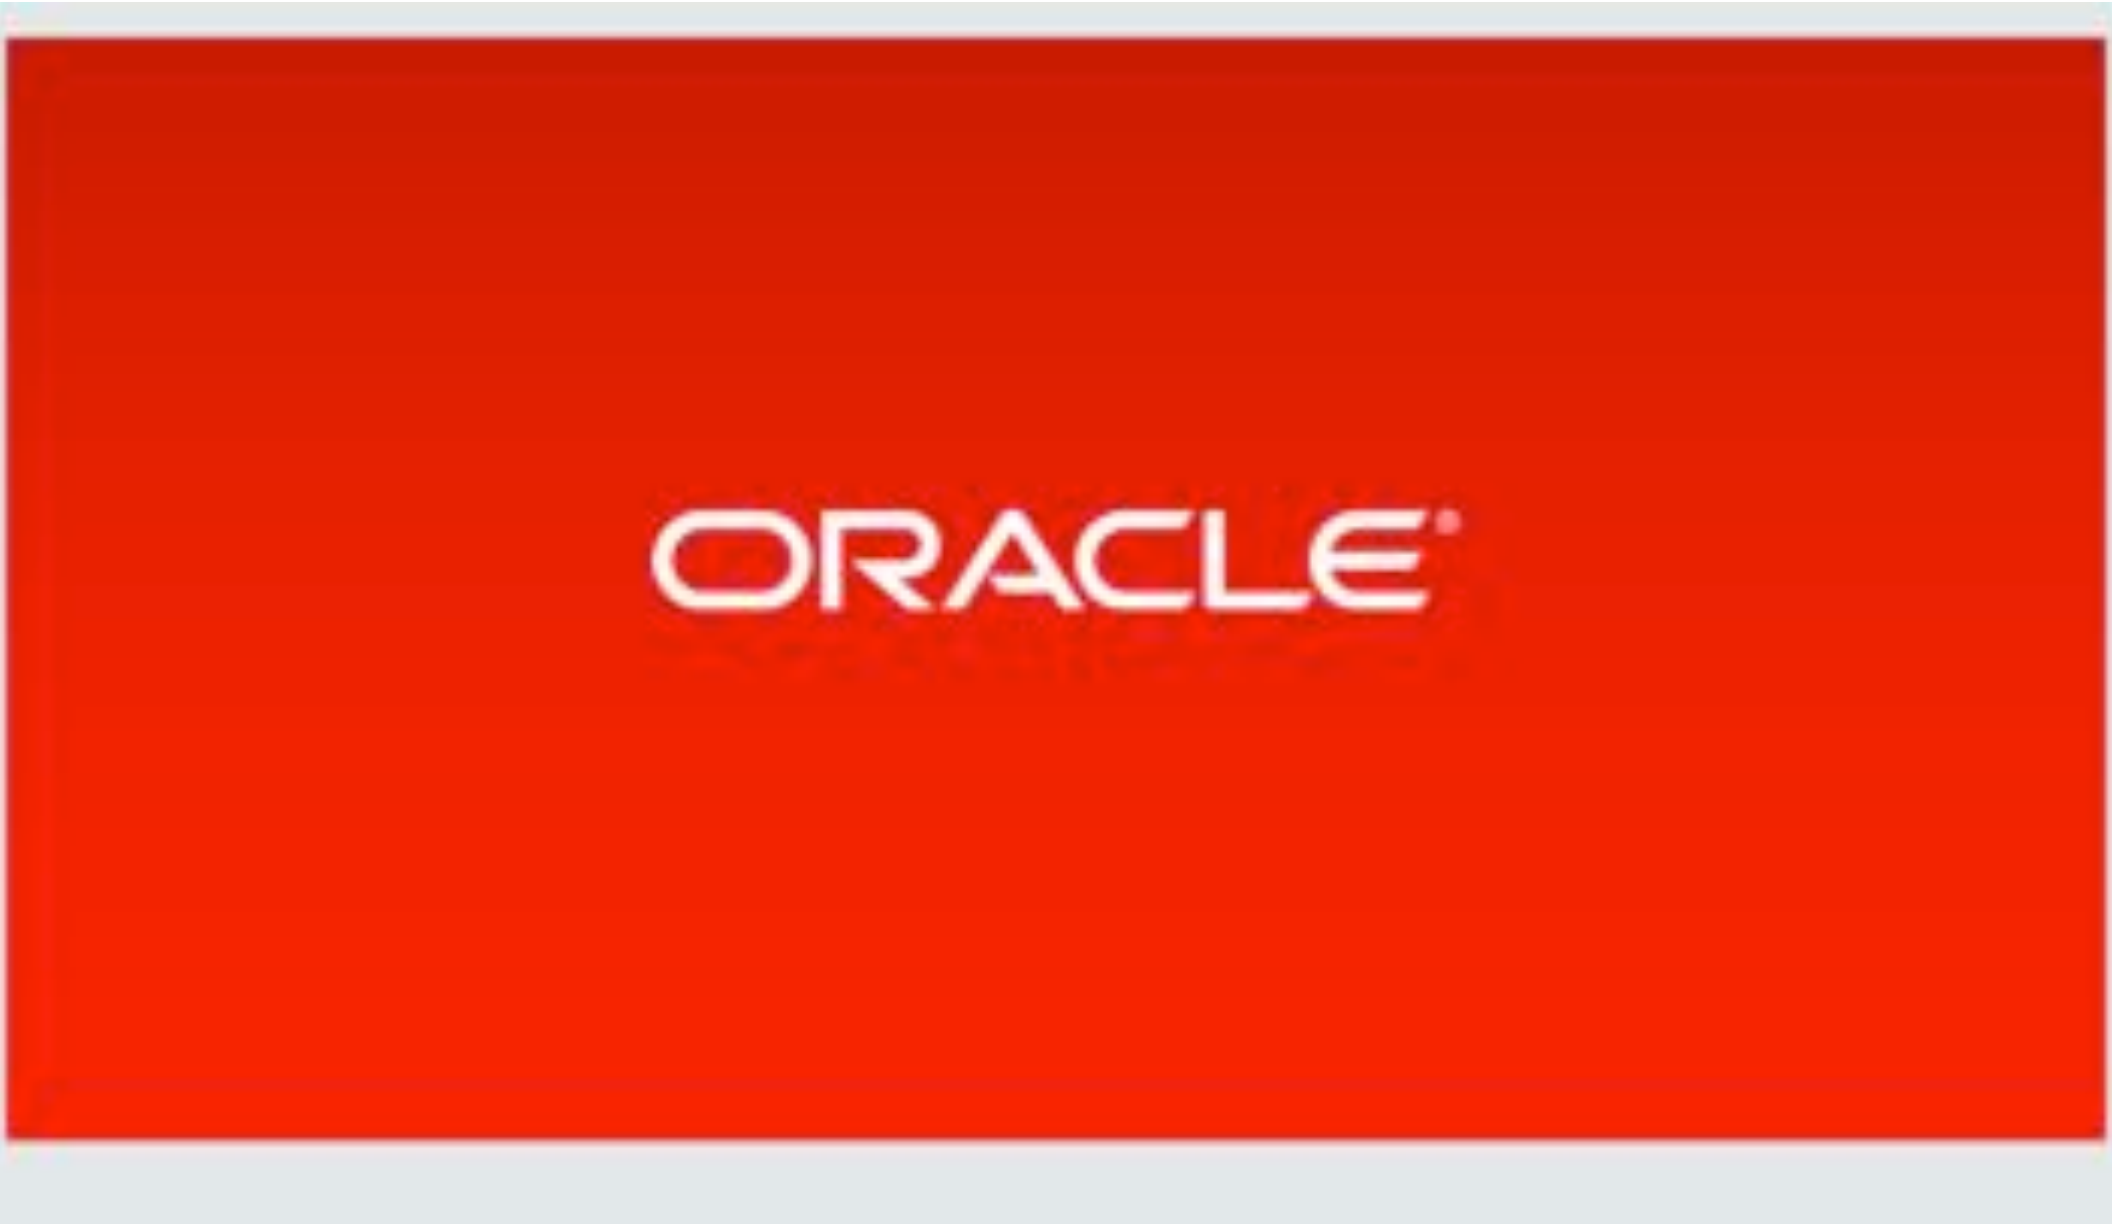
ORACLE