

ORACLE®

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.

ORACLE®
TimesTen

ORACLE®
TimesTen Scaleout

Oracle TimesTen In-Memory Database Release 18.1

Ananth Raghavan, VP Oracle TimesTen InMemory Database
Radu Parvu, Senior Data Architect, Accenture

ORACLE®

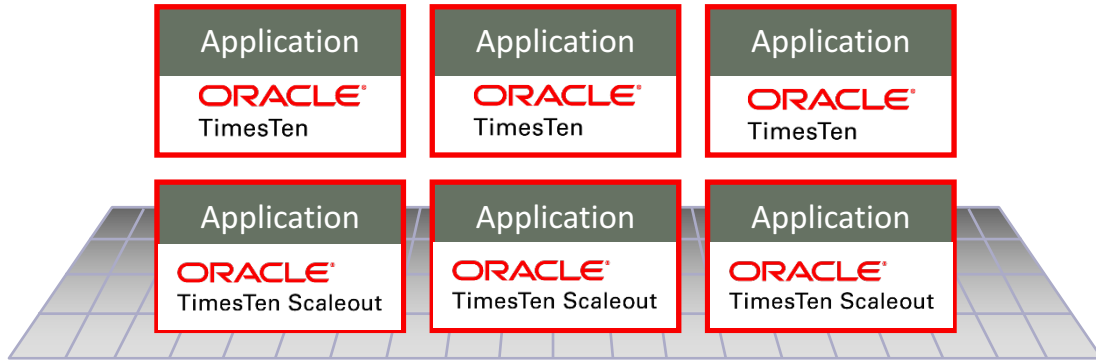
Agenda

- 1 Introduction
- 2 TimesTen Classic
- 3 Writing Applications for TimesTen
- 4 TimesTen Application-Tier Database Cache
- 5 TimesTen Scaleout


Introduction

Best In-Memory Databases: For Both OLTP and Analytics

In-Memory for OLTP



In-Memory for Analytics



Oracle TimesTen In-Memory Database

- Lightweight, highly-available IMDB
- Primary use case: **Extreme OLTP**
- **Microsecond** response time
- **Millions of TPS** on commodity hardware

Oracle Database In-Memory Option


- **Dual-Format In-Memory Database**
- Primary use case: **Real Time Analytics**
- **Billions of Rows/Sec** scan rate
- Faster mixed-workload enterprise OLTP
 - Fewer indexes needed to support analytics

The Forrester Wave™: In-Memory Databases, Q1 2017

**Oracle In-Memory Databases
Scored Highest by Forrester
on both Current Offering
and Strategy**

<http://www.oracle.com/us/corporate/analystreports/forrester-imdb-wave-2017-3616348.pdf>

The Forrester Wave™ is copyrighted by Forrester Research, Inc. Forrester and Forrester Wave™ are trademarks of Forrester Research, Inc. The Forrester Wave™ is a graphical representation of Forrester's call on a market and is plotted using a detailed spreadsheet with exposed scores, weightings, and comments. Forrester does not endorse any vendor, product, or service depicted in the Forrester Wave. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change.



Most Widely Used Relational In-Memory Database

Deployed by Thousands of Companies




TimesTen

Relational Database



- Pure in-memory
- ACID compliant
- Standard SQL
- Entire database in RAM

Persistent and Recoverable




- Database and Transaction logs persisted on local disk or flash storage
- Automatic recovery after failure

Extremely Fast



- Microseconds response time
- Very high throughput

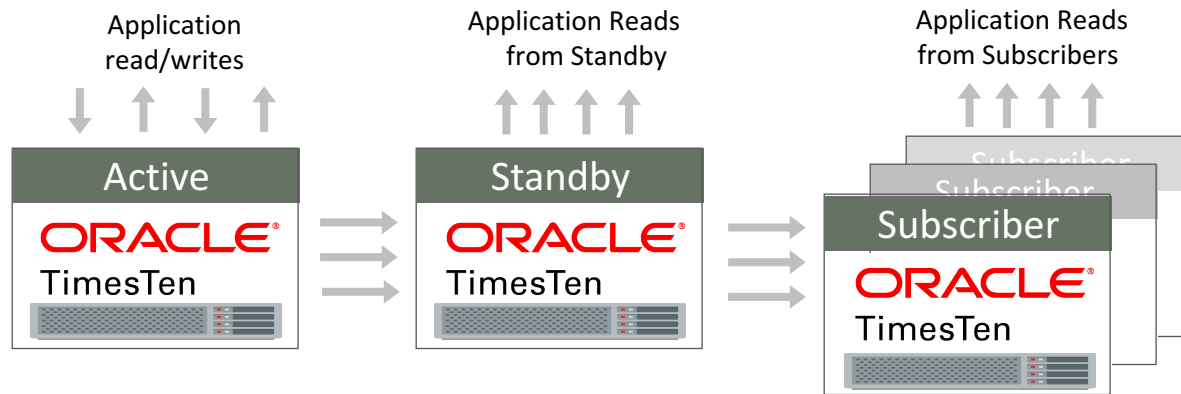
Highly Available



- Active-Standby and multi-master replication
- Very high performance parallel replication
- K-Safety
- HA and Disaster Recovery

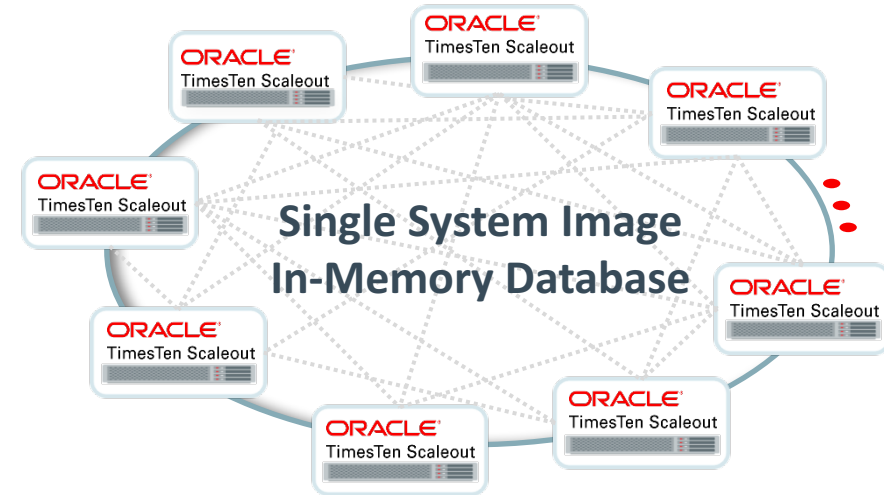
TimesTen In-Memory Database 18.1

TimesTen Classic



- Replicated In-Memory Relational Database
- Highly Available
- **Extremely low latency** reads and writes
- **Read scaling** across multiple hosts
- Cache functionality for Oracle Database

TimesTen Scaleout

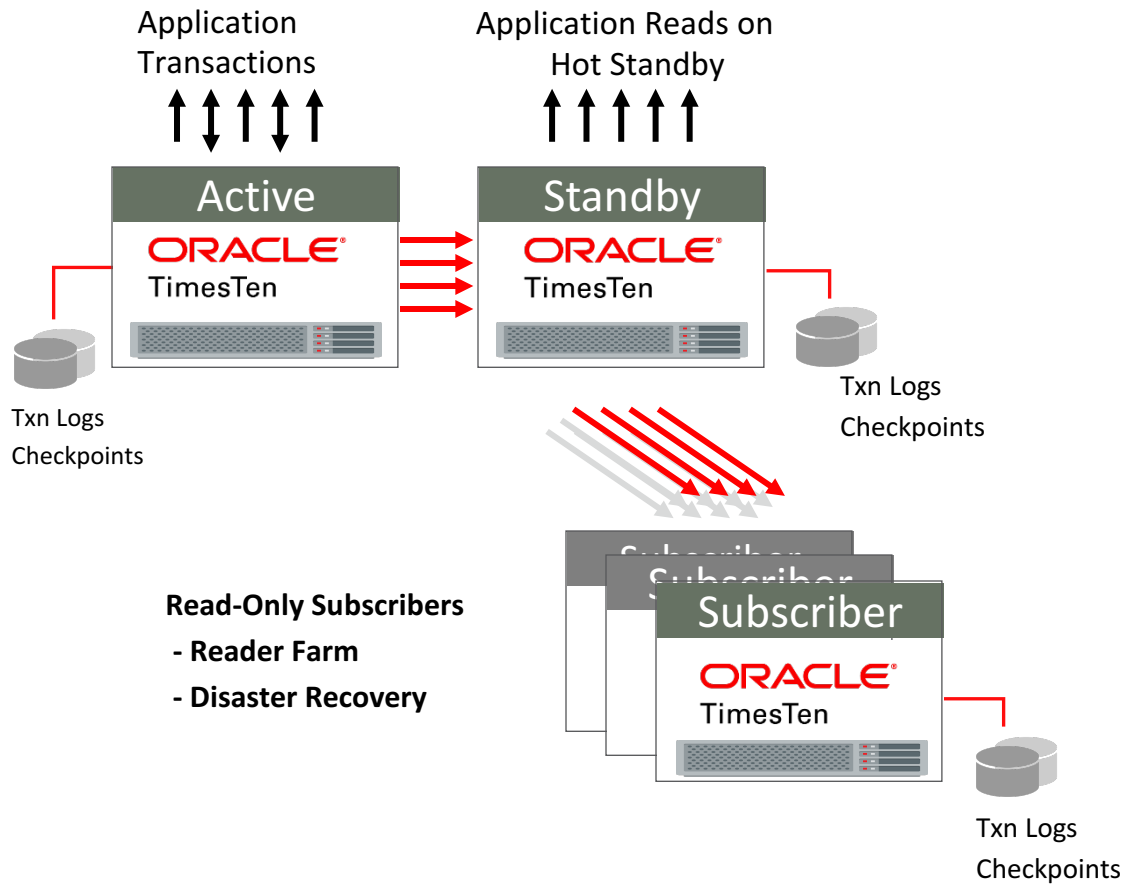


- Scale-Out In-Memory Relational Database
- Highly Available
- **Extremely high throughput** reads and writes
- **Scales both reads and writes**

TimesTen Classic

Real-Time Transactional Replication

High Availability and Disaster Recovery

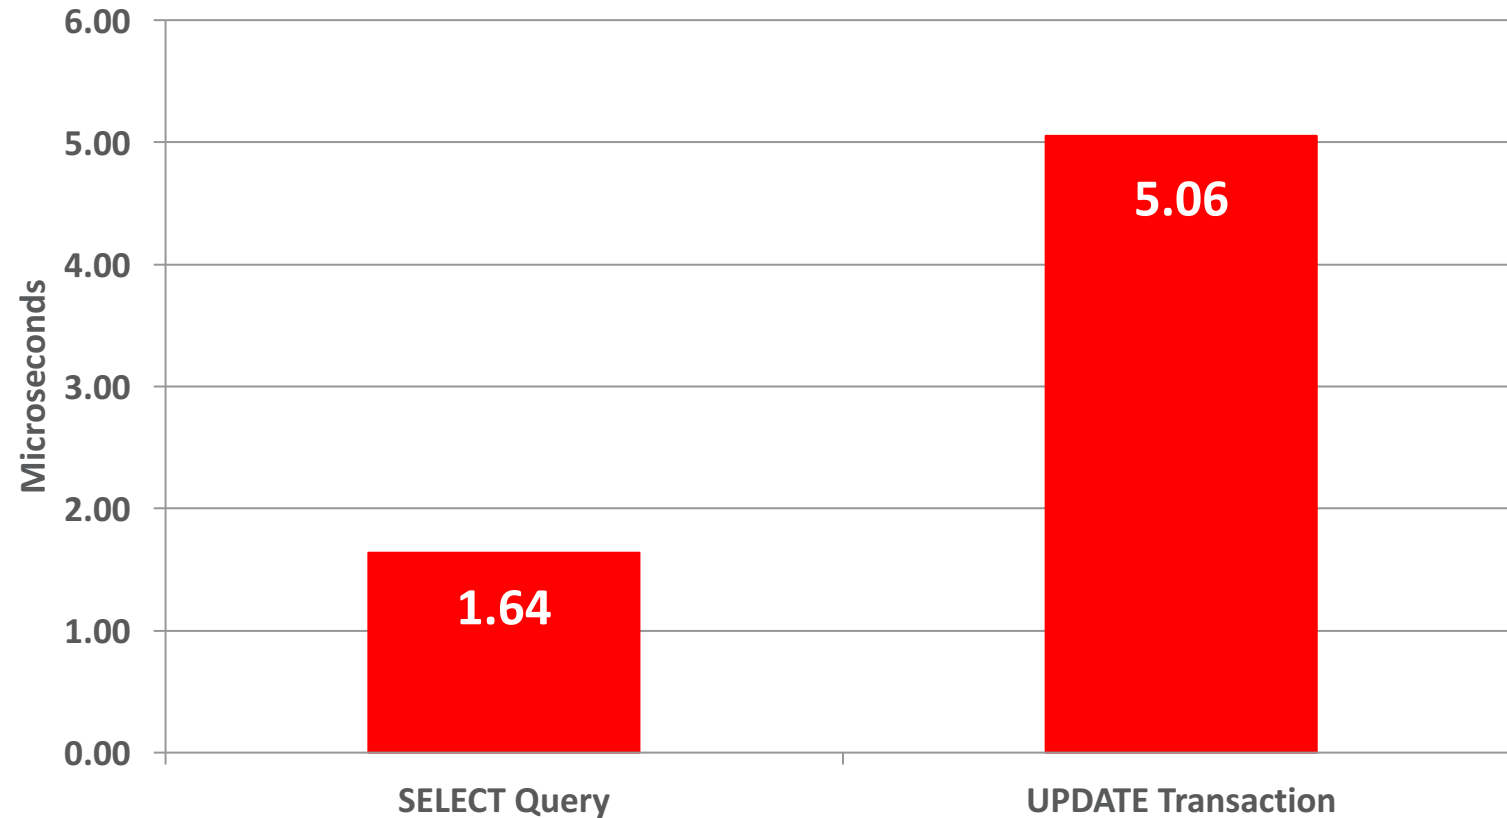


- High performance
 - Synchronous / Asynchronous
 - Parallel send of log streams
 - Parallel apply of changes on Standby and Subscribers
- HA and DR support
- Online rolling upgrades
 - No application downtime
 - Cross-version replication
- Integration with Oracle Clusterware

Performance – Response Time

Low Latency - **Microseconds** Response Time

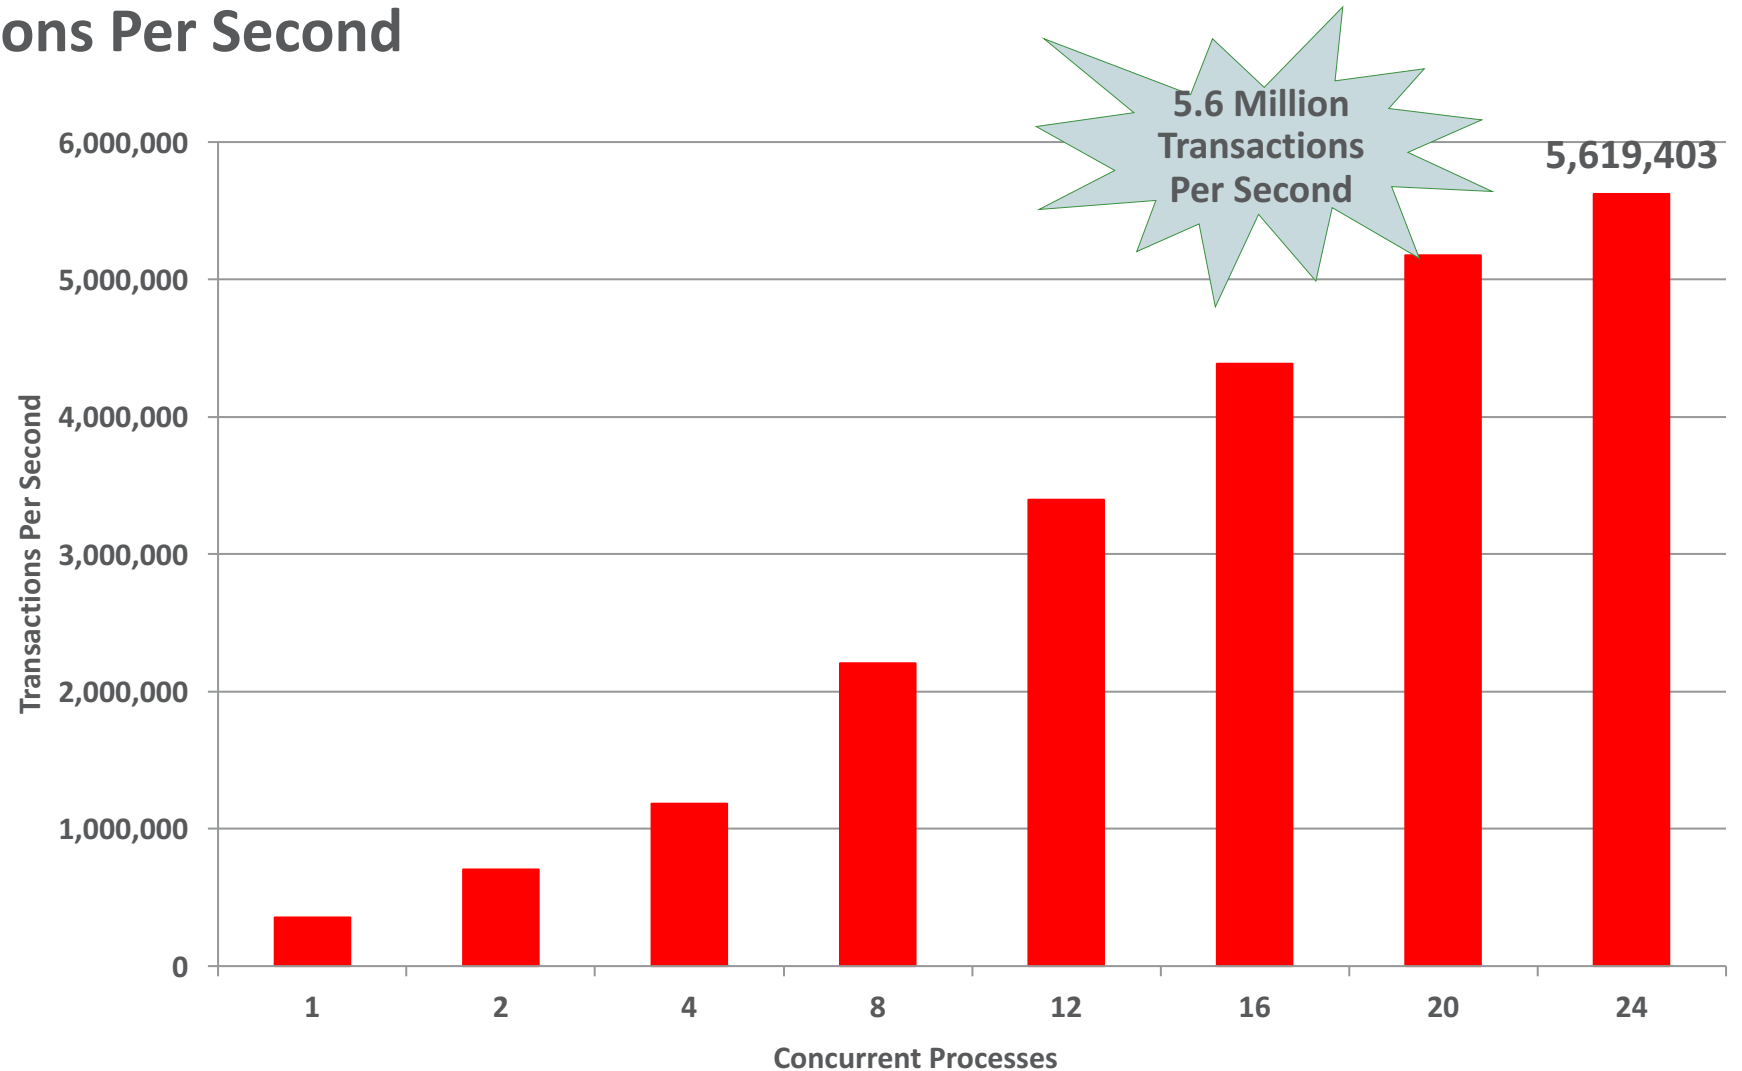
TPTBM Read and Update
E5-2699 v4 @ 2.20GHz
2 socket, 22 cores/socket,
2 threads/core
TimesTen 11.2.2.8.0
(100M rows, 17GB data)




Performance - Throughput

5.6 Million Transactions Per Second

TPTBM Mixed
Workload
(80%R-10%U-5%I-5%D)
E5-2699 v4 @ 2.20GHz
2 socket, 22
cores/socket,
2 threads/core
TimesTen 11.2.2.8.0
(100M rows, 17GB)



Application Development

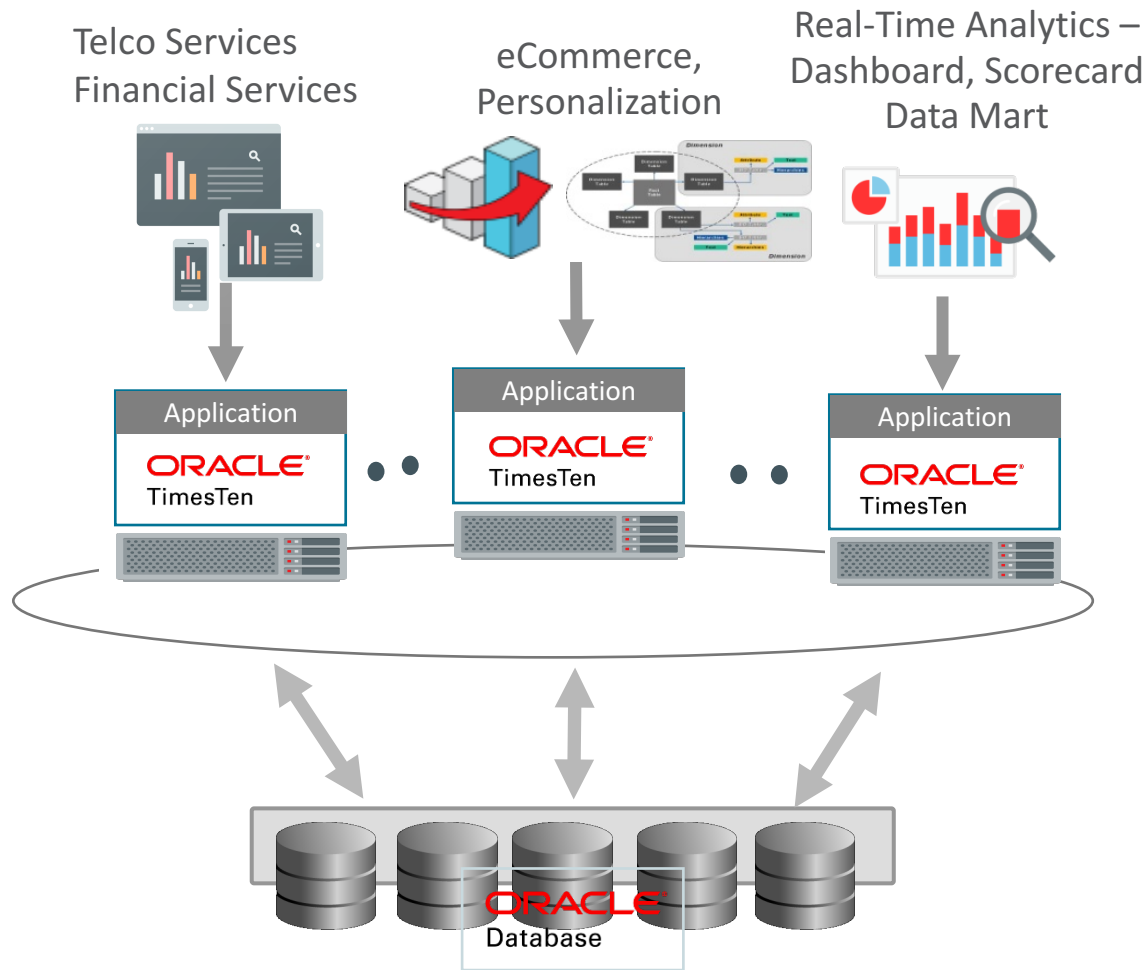


- Industry standard and Oracle Compatible APIs
- Flexible application development
- Nearly any programming language
- Nearly any environment

TimesTen Classic Summary

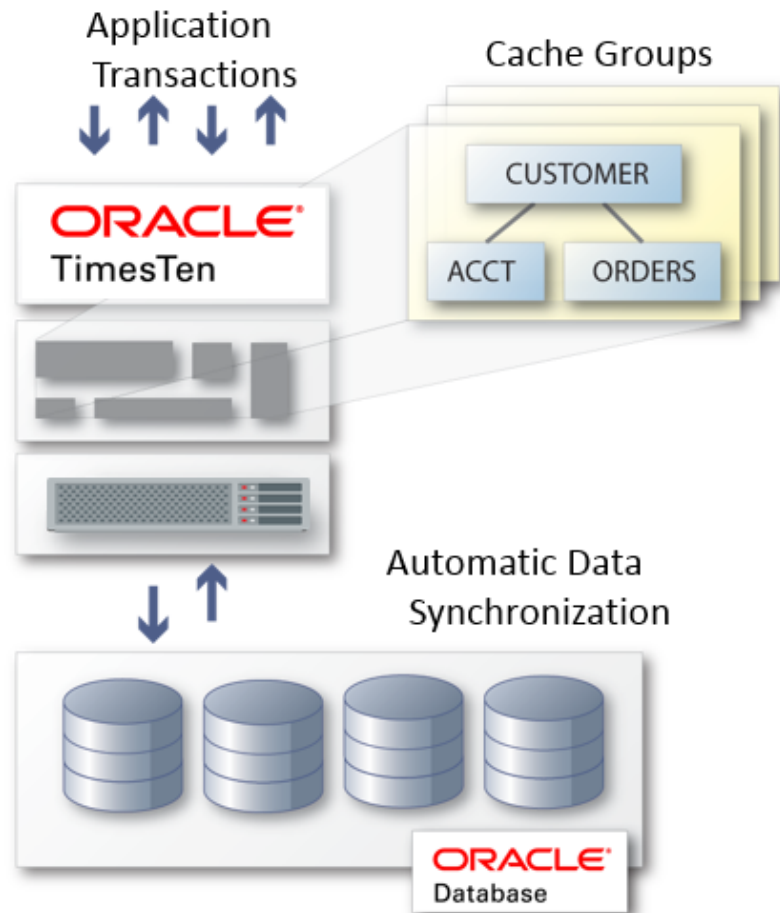
- Respond to real-time events
Response time measured in microseconds
- Provide consistent level of responsiveness
Fast and consistent response time with low latency
- Provide continuous, uninterrupted service
High availability and online upgrades
- Ability to leverage existing applications with minimal changes to application code and interfaces
Standard SQL, relational model, standard APIs
- Compatible with Oracle products
Oracle Enterprise Manager, SQL Developer, Oracle GoldenGate, Oracle Clusterware

TimesTen Application-Tier Database Cache for Oracle Database



- Cache subset of Oracle Database tables in TimesTen for better response time
 - With full persistence to local storage
- Read-write caching
 - Transaction execution and persistence in TimesTen
- Read-only caching
 - Transactions executed in Oracle Database
- Same architecture as TimesTen Classic
 - Supports cache tables and native TimesTen tables
- HA and fault tolerance in the application-tier

Flexible Cache Group Configurations

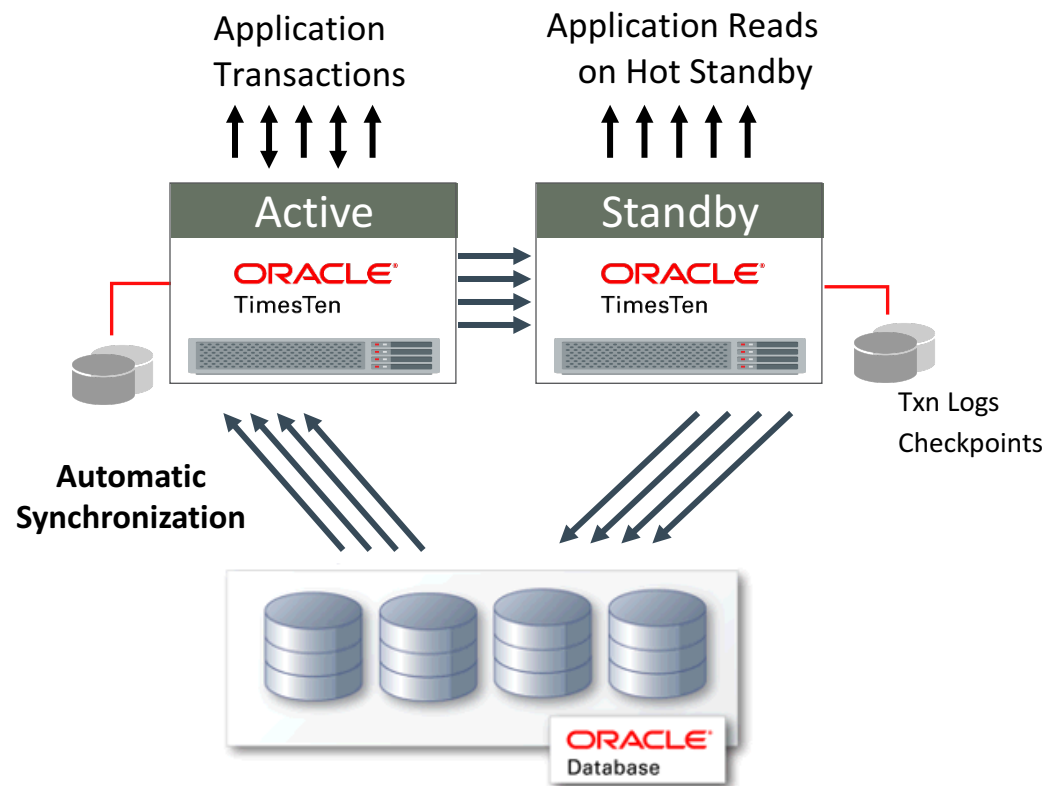


- Cache Group describes the Oracle Database tables to cache
 - All or subset of rows and columns
 - Defined using SQL

```
CREATE CACHE GROUP PremierUsers
FROM OE.CUSTOMER (
    NAME VARCHAR2(100) NOT NULL,
    ADDR VARCHAR2(100)
)
WHERE OE.CUSTOMER.ORDER > 500;
```
- Cache tables are regular tables in TimesTen
 - Joins/search, insert/update/delete

Read-write and Read-only Caching

High Availability - MAA



- Read-write caching
 - Parallel replication of transactions from Active to Standby
 - Parallel write-through of transactions to Oracle Database
- Read-only caching
 - Multi-stream refresh of transactions from Oracle Database
 - Parallel replication of refresh transactions to Standby
- Application continues even if Oracle Database connection is down

TimesTen Cache Summary

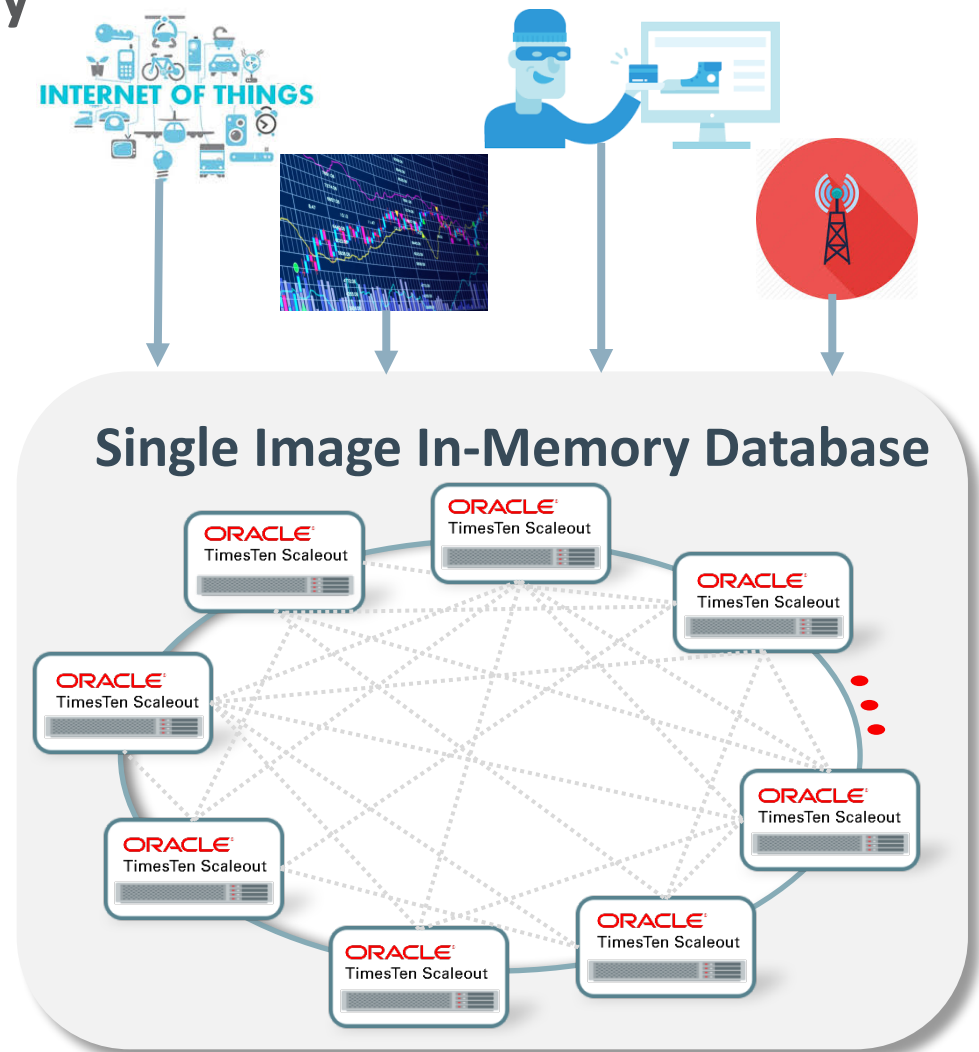
- Accelerating existing Oracle Database Applications
 - Caching from Oracle Database with automatic change synchronization*
- Same architecture and features as TimesTen Classic
 - Low, consistent response time*
 - High availability and online upgrades*
 - Standard SQL, relational model and standard APIs*
 - Compatible with Oracle Enterprise Manager, SQL Developer, GoldenGate and Clusterware*
- Multiple configuration options – mix and match
 - Read-only cache groups*
 - Write-through cache groups*
 - Native TimesTen tables*

TimesTen Scaleout

Distributed, Shared Nothing, In-Memory Database

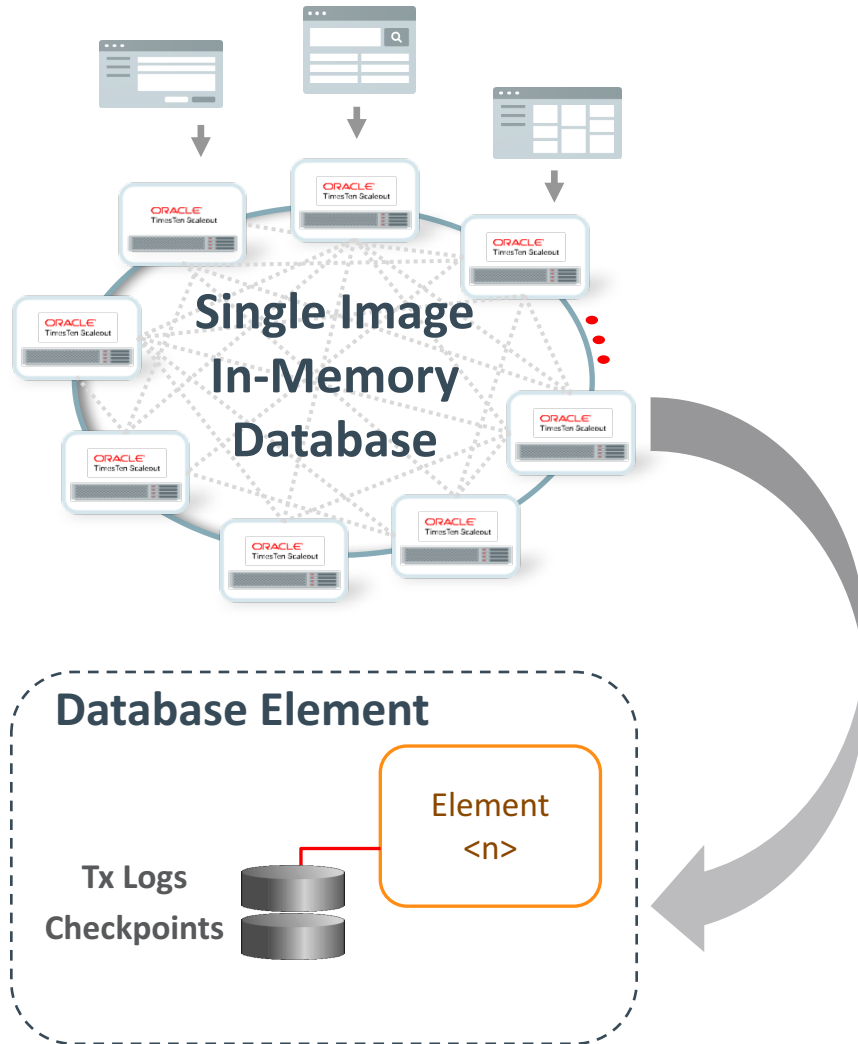
Single-Image DB with High Availability and Elasticity

- For High-Velocity **Extreme OLTP** applications
 - IOT, trading, fraud detection, mobile, click stream, billing, orders, etc.
- Appears as a single DB to applications
 - **Not** as a sharded database
- Adding and removing DB elements
 - Data automatically redistributed
 - Workload automatically uses new elements
- Built-in HA via fully-active element copies
 - Element copies automatically kept in sync
- Highly compatible with Oracle
 - Data types, APIs, SQL & PL/SQL



TimesTen Scaleout - Database Elements

Unit of Persistence and Recovery



- Each database consists of *elements*
- Each *element* stores a portion of data from its database
- Each *element* has its own set of checkpoint files and transaction log files for persistence
- The *element* is the smallest unit for database persistence, failure recovery and high availability

TimesTen Scaleout - Database Elements

A “logical” look

- Each element contains:
 - Information about all *users* in the database
 - The *schema* of the entire database
 - Some rows of each *table* in the database



Element 1

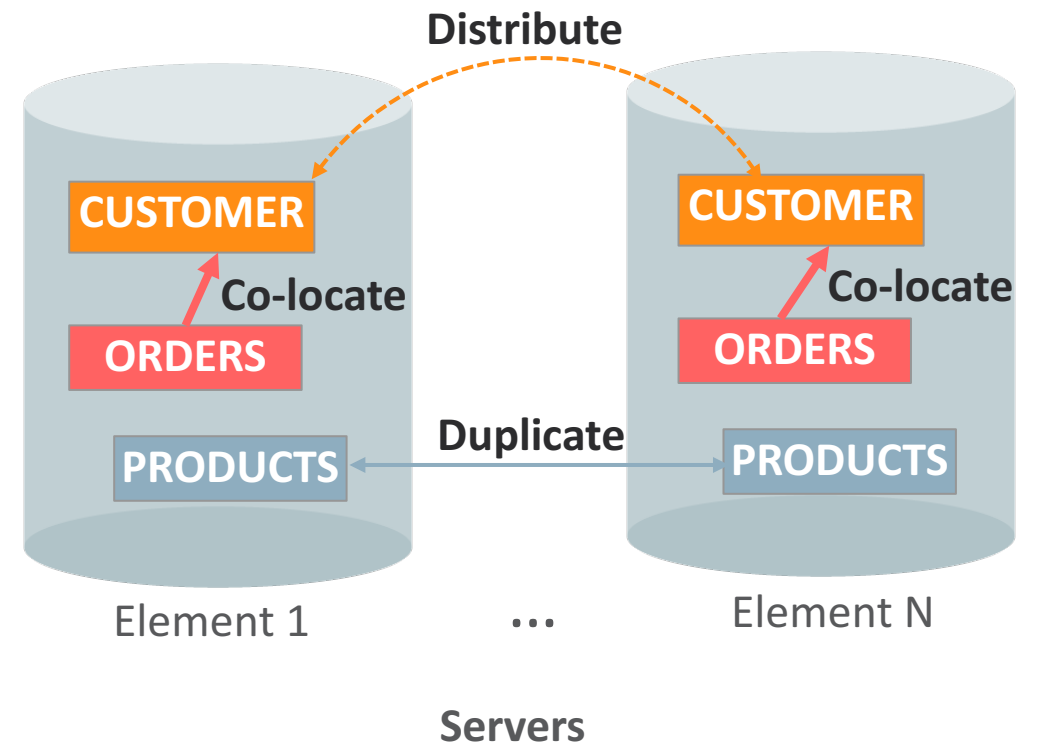


Element 2

TimesTen Scaleout - Data Distribution

Specified at the table level


- **DISTRIBUTE** large tables by consistent hash
 - Distribute CUSTOMER rows on all elements by hash of Customer ID
- **COLOCATE** child table rows with parent table row to maximize locality
 - Place ORDERS rows in same element along with corresponding CUSTOMER row
- **DUPLICATE** small read-mostly tables on all elements for maximum locality
 - Duplicate the PRODUCT list on all elements



Distribute by Hash

- Consistent hash algorithm
- By hashing the distribution key column(s) or primary key column(s)
- Rows are “randomly” and evenly distributed across elements
- The default distribution method
 - There are ‘K’ copies of each row for HA, where ‘K’ is the K-safety factor
- Appropriate for most tables

```
CREATE TABLE CUSTOMER (  
  ID NUMBER NOT NULL PRIMARY KEY,  
  NAME VARCHAR2(100)  
) ... DISTRIBUTE BY HASH;
```

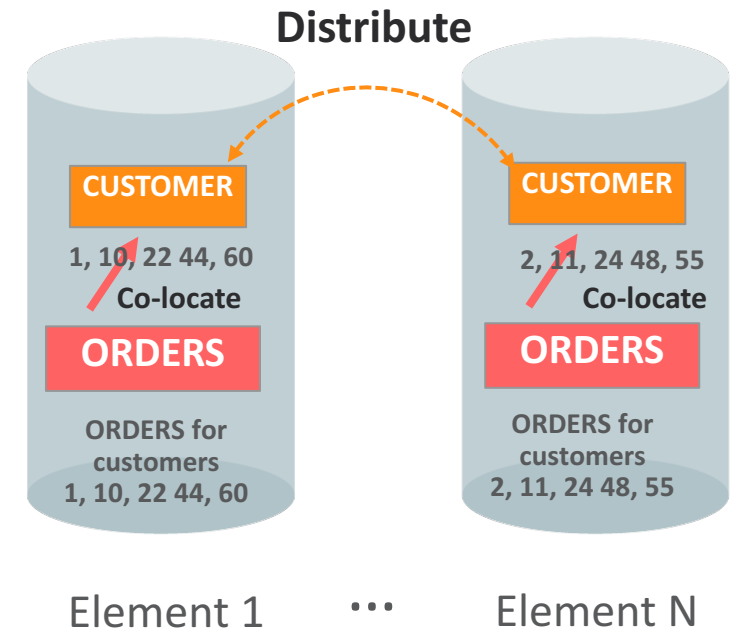


Distribute by Reference

- “Child” rows are located in the same elements as “parent” rows
- Foreign keys define “parents” and “children”
- Appropriate for tables that:
 - Are logically “children” of a single “parent” table
 - Parent and child will often be referenced together in queries
- Locating related data together provides best performance provided access is mainly via the ‘reference’ FK

```
CREATE TABLE CUSTOMER (  
  ID NUMBER NOT NULL  
  PRIMARY KEY,  
  NAME VARCHAR2(100)  
  ...  
) DISTRIBUTE BY HASH;
```

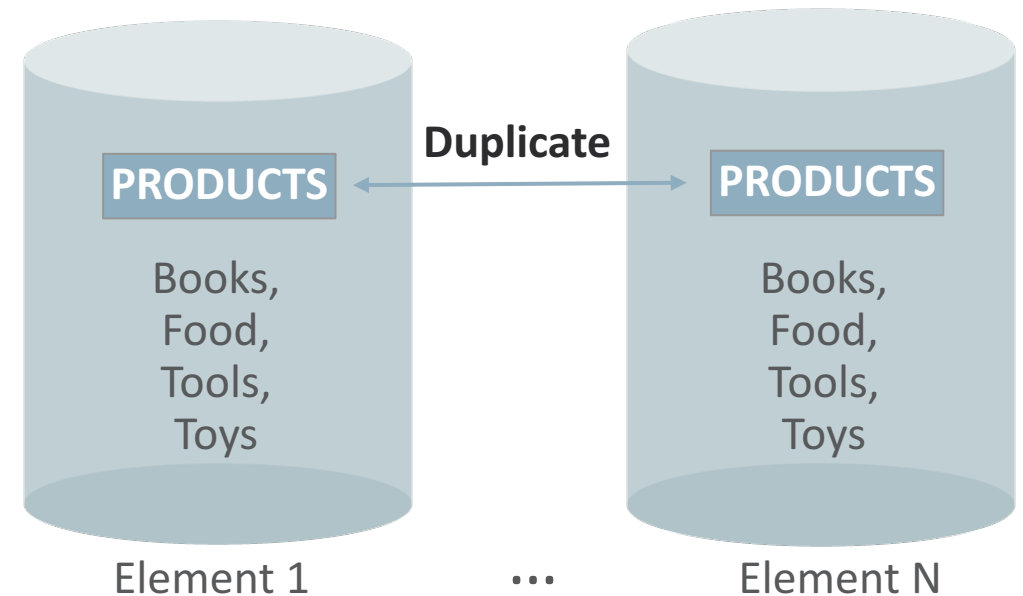
```
CREATE TABLE ORDERS (  
  ID NUMBER NOT NULL  
  PRIMARY KEY,  
  CUST_ID NUMBER NOT NULL,  
  FOREIGN KEY (C)  
  REFERENCES CUSTOMER(ID),  
  ...  
) DISTRIBUTE BY REFERENCE;
```



Duplicate


- Every row is present in every element of the grid
- Appropriate for tables that are:
 - Relatively small
 - Frequently read
 - Infrequently modified

```
CREATE TABLE PRODUCTS (  
  Prod_ID      NUMBER  
              NOT NULL  
              PRIMARY KEY,  
  Prod_Name    CHAR(12)  
) DUPLICATE;
```



TimesTen Scaleout - High Availability

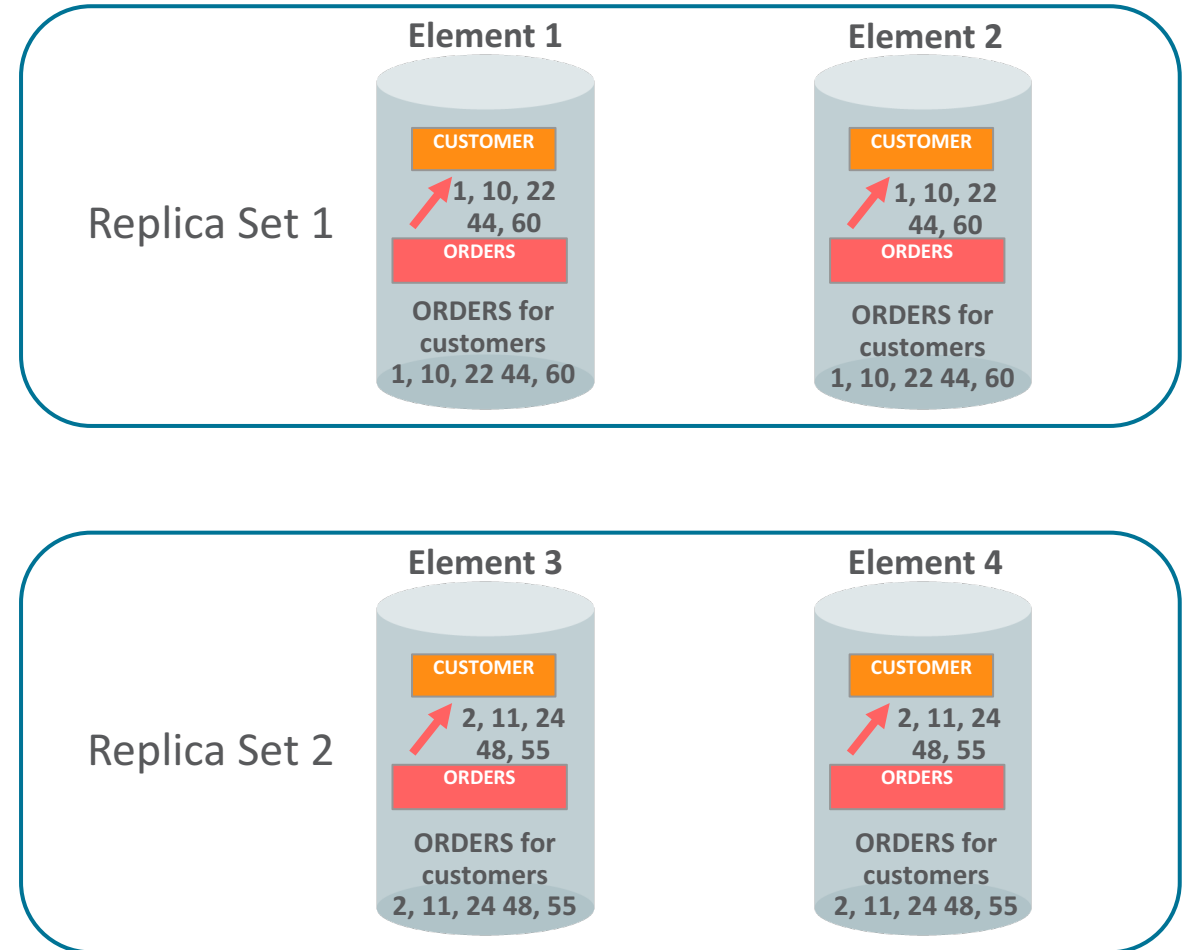
K-safety, All Active



- Built-in HA via multiple copies of the data (K-safety)
 - Automatically kept in sync
- **All** replicas are **active** for **reads** and **writes**
 - Increases the compute capacity
- Transactions can be initiated from, and executed on, any replica
- Queries and transactions can span any/all elements

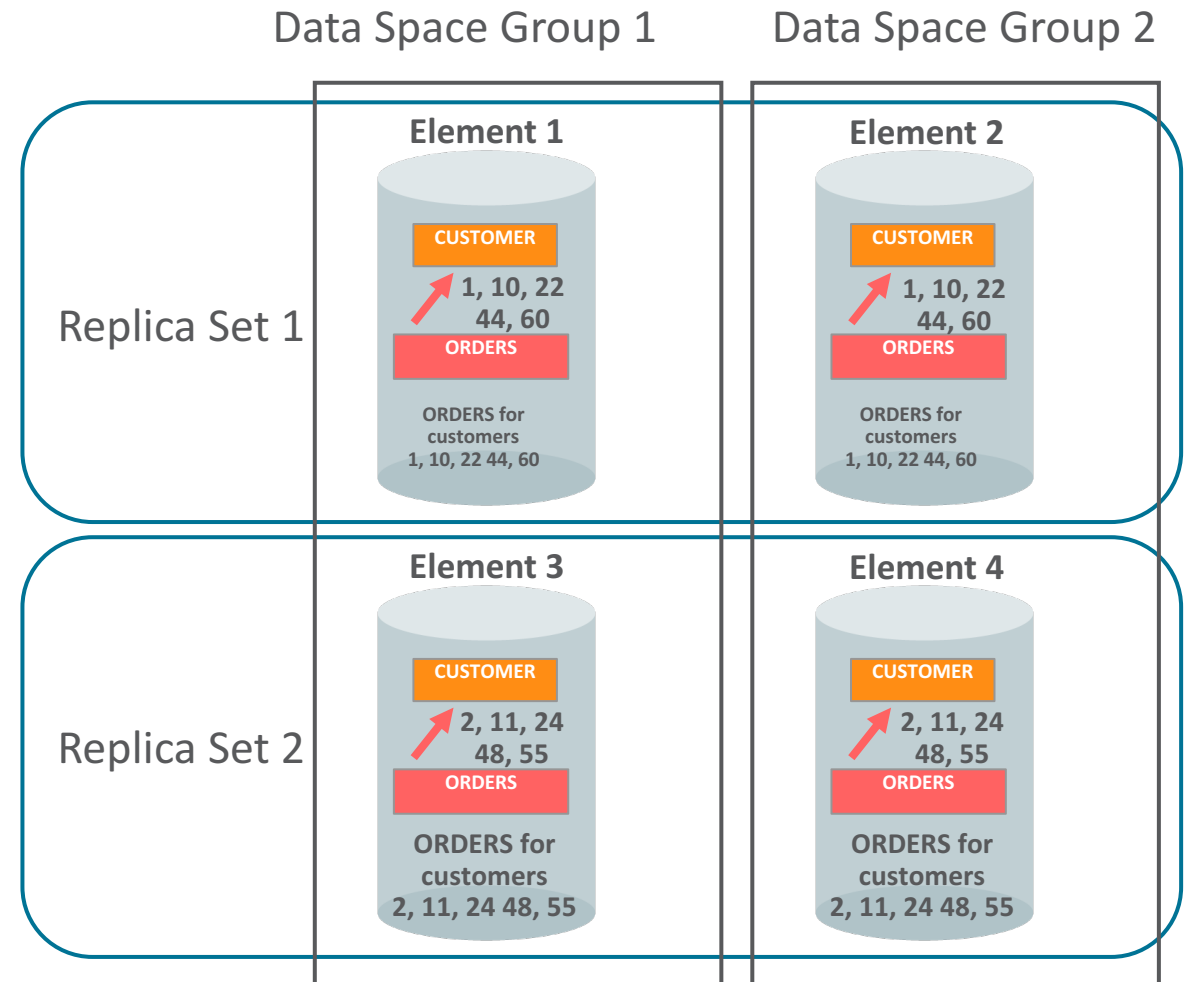
TimesTen Scaleout - Replica Sets

- Elements of a database are logically grouped into *replica sets*
- Each *replica set* contains K elements
- Elements in a *replica set* contain exactly the same data
- All elements in a *replica set* are “active”
 - Two phase commit protocols keep them in sync
- *Replica sets* are automatically created and managed



TimesTen Scaleout - Data Space Groups

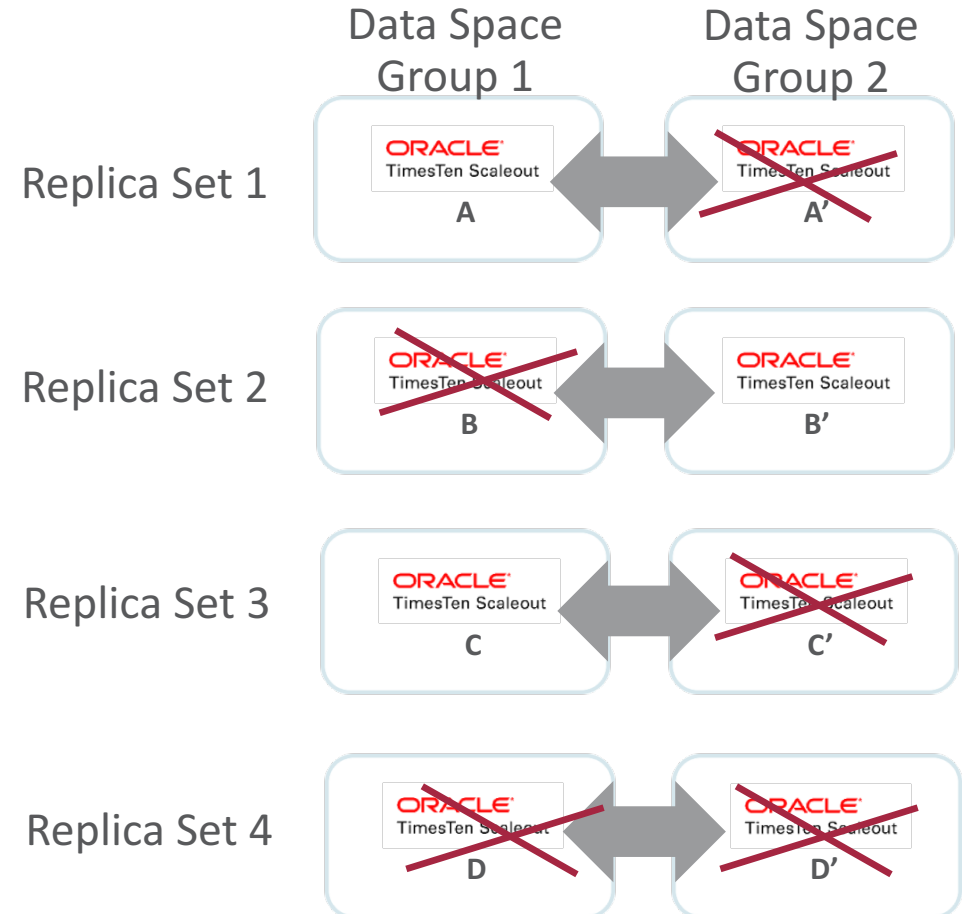
- When defining a Grid, the instance administrator must assign each host to a *data space group*
- If $K > 1$, each host must be assigned to a specific data space group
- Each replica set is automatically created from one element on a host in each data space group
- Maximizes data availability by minimizing common failure points



Database Fault Tolerance – No Application Down Time

Provided one full copy of the database is available

- If multiple elements fail, applications will continue provided there is one complete copy of the database
- Elements recover automatically after failure
- If an entire replica set is down, that data is unavailable until it recovers
 - Application can **explicitly** choose to accept partial results

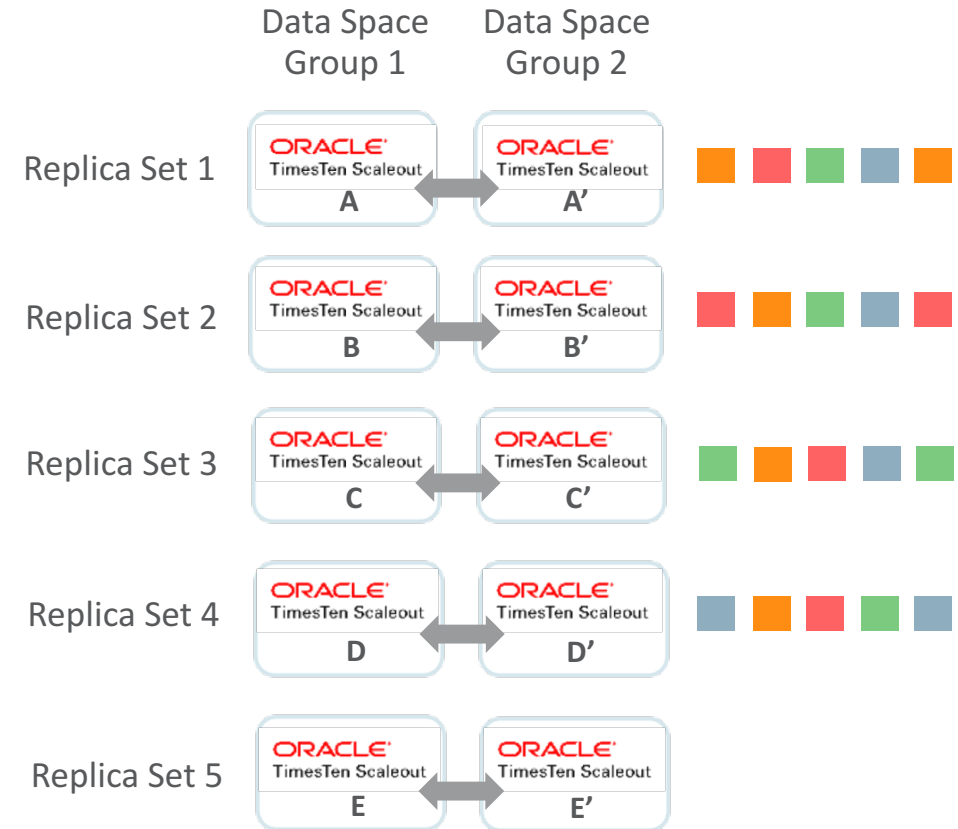


TimesTen Scaleout - Elastic Scalability

Expand and shrink the database based on business needs

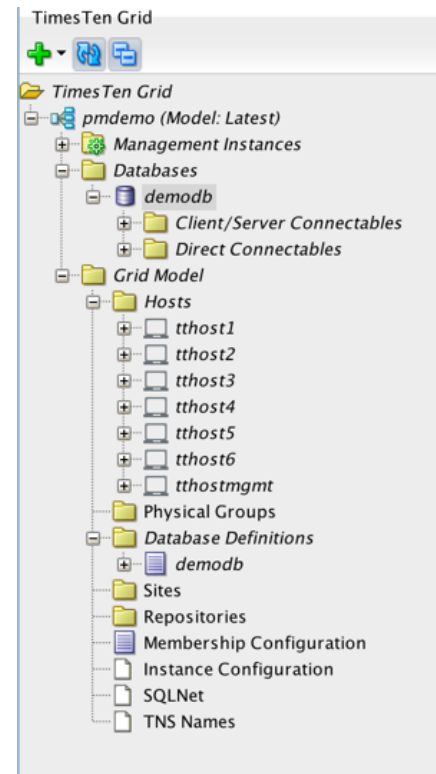
Adding (and removing) database elements

- Data redistributed to new elements
- Workload automatically uses the new elements
- Connections will start to use new elements
- Throughput increases due to increased compute resources



Centralized Installation and Management

- All TimesTen Scaleout management and admin operations can be performed from a single host
 - Installing software
 - Patching software
 - Configuration
 - Database creation and management
 - Backup and restore
 - Monitoring
 - Collecting diagnostics
- Command line and SQL Developer UI interfaces



Status Database Definition Topology

Database demodb status is: created, loaded-complete, open

Number of application connections to demodb: [0](#)

Number of system connections to demodb: 168

Database distributed in 6 instances

| Element ID | Host name | Instance Name | In Distribution Map | Data Space Group |
|------------|-----------|---------------|---------------------|------------------|
| 1 | tthost1 | instance1 | Yes | 1 |
| 2 | tthost2 | instance2 | Yes | 2 |
| 3 | tthost3 | instance3 | Yes | 1 |
| 4 | tthost4 | instance4 | Yes | 2 |
| 5 | tthost5 | instance5 | Yes | 1 |
| 6 | tthost6 | instance6 | Yes | 2 |

China Mobile Marketing Promotion System

Chongqing Mobile Subsidiary

Application Overview

- Industry : Telecom
- Business : Business & Operation Support System
- Application : Marketing Promotion System
 - Promote China Mobile products to various channels including website, APPS, SMS, WeChat, etc.
 - 30 million target subscribers
 - 15 million successful promotions per day

Challenges

- Highly concurrent mobile locations based query
 - For every subscriber in the mobile carrier network
- High transaction throughput with consistent low latency
 - Read mostly application with small amount of DML and DDL
- Scalability to achieve higher throughput

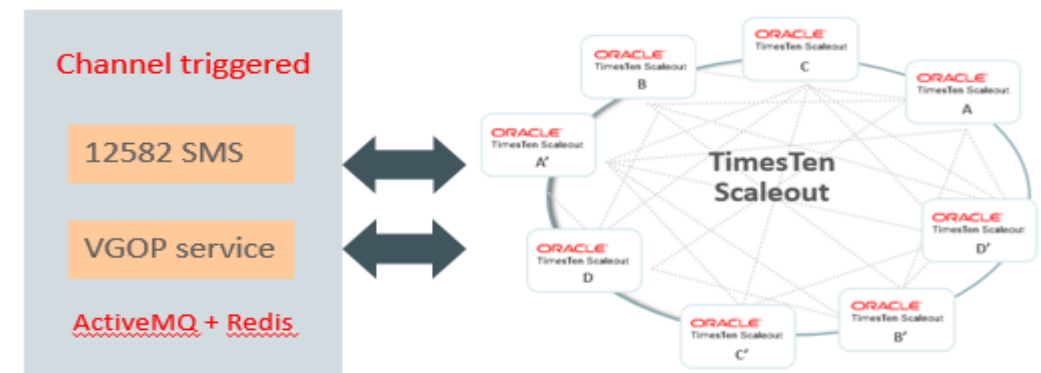
Solution

- TimesTen Scaleout with K=2 for High Availability

Why TimesTen Scaleout ?

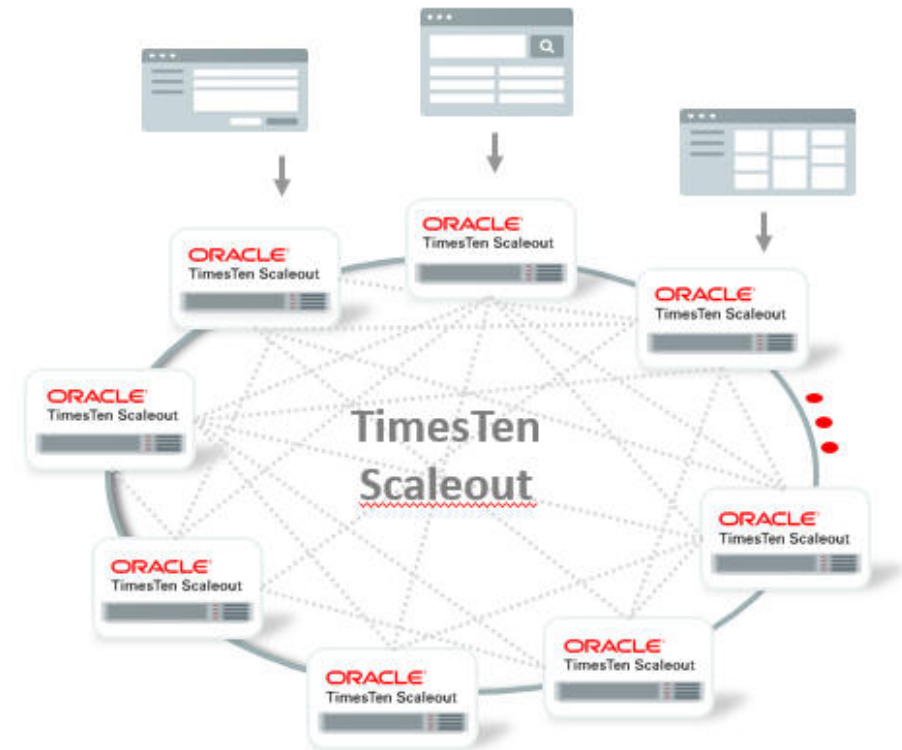
- End-to-end response time ~200 milliseconds via C/S connection mode
- New LBS (location based service) module with 2000 concurrent connections in peak time
- Easy, automatic high-availability
- No application code changes moving from TimesTen 11.2.2 to TimesTen Scaleout
- Scalability for future growth

Marketing Promotion System currently supports over 30 million subscribers and populate 15 million marketing messages per day.



TimesTen Scaleout Summary

- Extreme performance
- Single database image, data location transparency
- Full SQL, ACID transactions
- Automatic high availability via K-safety
- Elastic scale-out and scale-in
- Easy to deploy and manage
- Easy application development
- On-premises or Cloud deployment



TimesTen Scaleout Performance

Running TimesTen Scaleout on Oracle Cloud Infrastructure

The **World's Fastest Cloud** [Network, CPU & Storage]
enables the **Worlds Fastest OLTP Database** (Oracle TimesTen Scaleout)
to achieve **extreme** Transaction Processing throughput

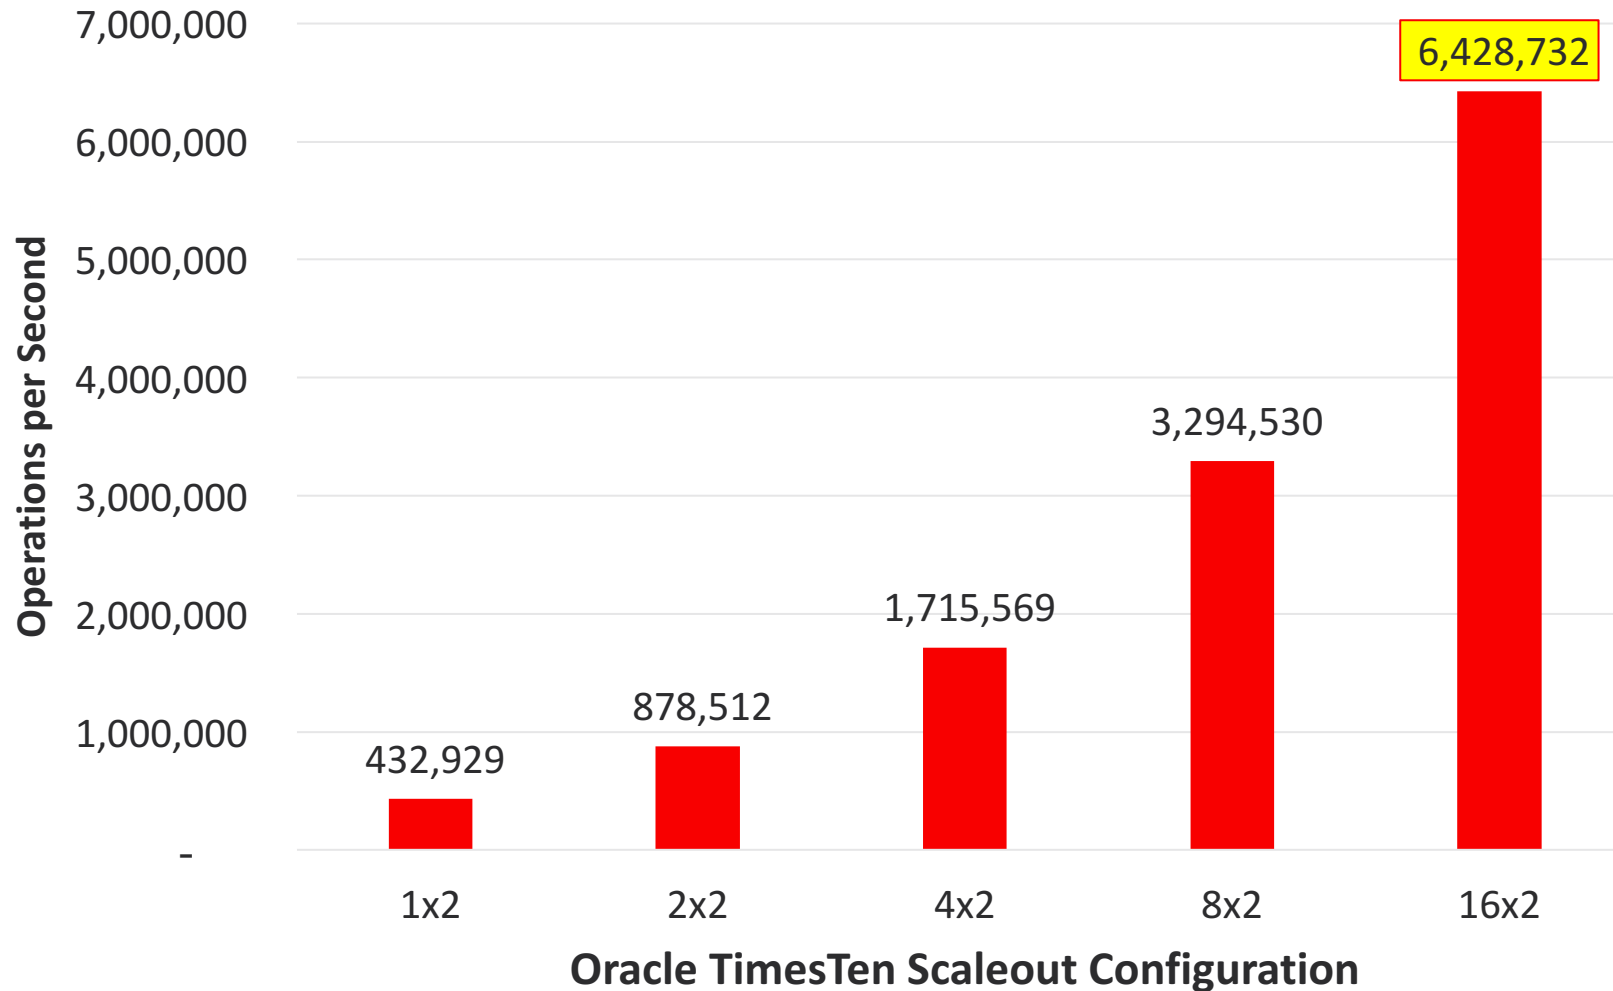
See what you can do with 32 **BM.DenseIO2.52** compute nodes!



What is the YCSB Workload?

- YCSB : **Y**ahoo **C**loud **S**erving **B**enchmark
 - Developed at Yahoo for Cloud Scale workloads
 - Widely used to compare scale-out databases, NoSQL databases, and (non-durable) in-memory data grids
- A series of workload types are defined:
 - Workload A: 50% reads, 50% Updates
 - Workload B: 95% reads, 5% Updates
 - Workload C: 100% reads
- The YCSB Client cannot be changed
 - DB Vendors implement the DB Client interface in Java
 - The version and exact configuration matters

YCSB Workload A (50% Read 50% Update): **6.4 Million Ops/Sec**



YCSB version 0.15.0

- 1KB record (100-byte x 10 Fields)
- 100M records / Replica Set
- Uniform Distribution

TimesTen Scaleout


- 1 to 16 replica sets
- 2 synchronous replicas per replica set

Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52

YCSB Workload B (95% Read 5% Update): **38 Million Ops/Sec**

Reminder: The best YCSB-B result found in our survey was 1.6 Million Ops/Sec



YCSB version 0.15.0

- 1KB record (100-byte x 10 Fields)
- 100M records / Replica Set
- Uniform Distribution


TimesTen Scaleout

- 1 to 16 replica sets
- 2 synchronous replicas per replica set

Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52

YCSB Workload C (100% Read): **65 Million Reads/Sec**



YCSB version 0.15.0

- 1KB record (100-byte x 10 Fields)
- 100M records / Replica Set
- Uniform Distribution

TimesTen Scaleout

- 1 to 16 replica sets
- 2 synchronous replicas per replica set

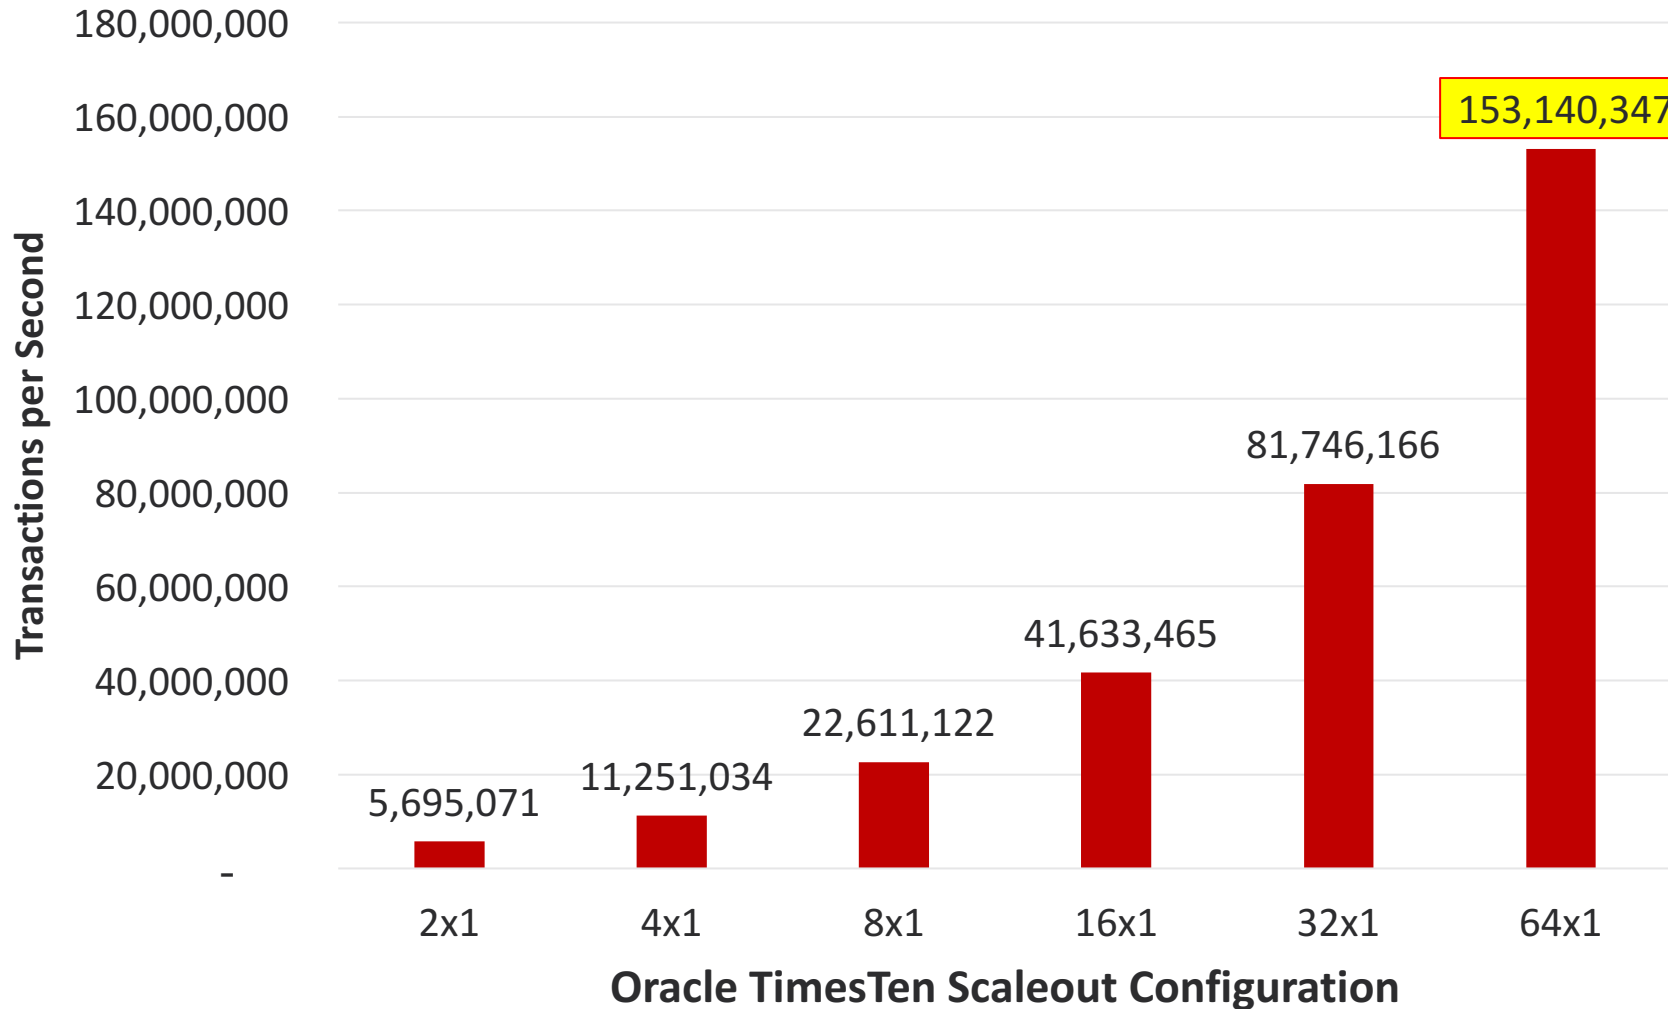
Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52

What is the TPTBM Workload?

- TPTBM : **T**elecom **P**rovider **T**hroughput **B**ench**M**ark
 - A benchmark originally developed by the TimesTen team
 - Represents common operations on a Telecom Subscriber database
 - Uses *standard* SQL and *standard* database APIs
 - Shipped with Oracle TimesTen as C and Java source code for the past 15 years
 - Quickly demonstrates the performance of user's hardware
- Common workload mixes:
 - 80% Reads, 20% Updates
 - 100% Reads
- The version and exact configuration matters

TPTBM 80% Read 20% Update: **153 Million Transactions/Sec**



TPTBM Configuration

- 128-byte record
- 100M records / Replica Set
- Uniform Distribution


TimesTen Scaleout

- 1 to 64 replica sets
- 1 replica per replica set

Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52
- Two TimesTen instances per compute node

TPTBM 100% Read: **1.4 Billion Reads Per Second!!**



TPTBM Configuration

- 128-byte record
- 100M records / Replica Set
- Uniform Distribution

TimesTen Scaleout

- 1 to 32 replica sets
- 2 synchronous replicas per replica set

Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52
- Two TimesTen instances per compute node

TimesTen Futures

- TimesTen 18.1 is ready for production use today!
- $K > 2$, i.e. We will support 3 or more copies of the data for even higher availability
- Online Scaleout and Scalein. The Grid can be elastically reshaped without any effect of running workloads
- Disaster Recovery: Support asynchronous replication to another TimesTen grid located in a different region.
- Geodistribution: Support a grid that spans regions across the globe.

TimesTen Futures

- Persistent Memory
 - Vastly Improved database startup time: less than 1 second vs more that 10 minutes
 - Much higher throughput for durable workloads: 6.5x
 - See a live demo at the Intel session on Tuesday 12:30 at YBCA
- Cloud Release
 - Scaleout Rollout Utility coming very soon. Will enable provisioning a Timesten Grid in less than 15 min.
 - Working on a fully managed Timesten cloud service


More TimesTen Sessions!

- Oracle TimesTen Scaleout: Getting Started [PR04015]
10/23 Room 3001, 11:15
- Intel Executive Session
YBCA, 10/23, 12:30
- Oracle TimesTen Scaleout: Developing Applications for OLTP, IoT and Machine Learning [PR04017], 10/24, 11:15
- Next-Generation OLTP: Oracle TimesTen Scaleout[HOL6319]
Marriott Marquis (Yerba Buena Level) - Salon 3/4
10/23, 5:15
10/24, 12:45

More TimesTen Sessions!

- Oracle Code One: How to Use the new JDBC 4.3 Sharding API for Massive OLTP Scaling [DEV4712]
Moscone West - Room 2003, 10/22, 4:00
- Oracle Code One: Database-Driven Machine Learning [DEV4714]
Moscone West - Room 2018, 10/25, 12:00
- Come see us at the Demogrounds!
 - Moscone South

**ORACLE
TIMESTEN
SCALEOUT: A
CLOUD-SCALE
IN-MEMORY
DATABASE
FOR OLTP
RADU PÂRVU -
SENIOR DATA
ARCHITECT-
ACCENTURE ENKITEC
GROUP**



accent[>]uretechnology

ACCENTURE ENKITEC GROUP

THE DATA DOCTORS



ELITE

- Global systems integrator focused on the Oracle platform
- Consultants average 15+ years of Oracle experience
- Worldwide leader in Exadata implementations
- 16 Oracle ACE members



EXPERTISE

ORACLE SPECIALIZATIONS

- Oracle Cloud IaaS
- Oracle Exadata
- Oracle Database
- Oracle Data Integrator & Golden Gate
- Oracle Data Warehouse
- Oracle Real Application Cluster
- Oracle Performance Tuning
- Oracle Database Security



SUCCESS




THOUGHT LEADERSHIP

Our consultants have been published in multiple subject areas and additional online resources that demonstrate Accenture's experience and expertise with the OES platform.



INNOVATION CENTER



RADU PÂRVU

SENIOR PRINCIPAL – ACCENTURE ENKITEC GROUP

Radu has over 18 years of experience working on database technology (mainly Oracle but other DBMS, too).

Radu holds License (one degree over Bachelor) of Engineering in Technological Physics from University of Bucharest.

Areas of Expertise:

Database: Data migrations, DB upgrades, backups and recovery, cloningData

Replication: Golden Gate mainly

In Memory DB Technologies

Business Intelligence and Data Warehousing:

ES / Exadata

Infrastructure Architecture: Maximum Availability Architecture and Disaster Recovery

CLIENTS NEED EXTREME INFRASTRUCTURE PERFORMANCE

- **Clients want everything:**
 - Extreme performance AND
 - Lower total IT costs
- **”Future” technologies are required today!**
 - Cloud
 - In-Memory
 - Artificial Intelligence
- **Our clients want:**
 - Sub-second response times under all circumstances
 - Ability to scale vertically and/or horizontally
- **TimesTen Scaleout is an essential building block for any future design that needs to support extreme OLTP workloads**

THE ARCHITECTS TOOLKIT

- **TimesTen Scaleout is an important component**
- **Oracle Database In-Memory feature has proven itself**
 - For DW and DSS workloads
 - Does not help for OLTP workloads
- **Previous OLTP options:**
 - TimesTen Classic: limited database size to single machine's RAM, or
 - Niche products with limited support / feature set
- **Until now there was no ACID compliant in-memory DBMS that would support large datasets!**

TIMESTEN SCALEOUT – OUR EXPERIENCE WITH IT

- **All our simulated workloads pass our performance test**
 - 50% Read – 50% Write
 - 95% Read - 5% Write
 - 100% Read
- **TimesTen Scaleout is orders of magnitude better when compared to any traditional RDBMS**
- **TimesTen Classic is even faster than TimesTen Scaleout**
 - Performance difference is noticeable for extreme workloads
- **Production ready:**
 - Remarkable stability for such a new, revolutionary architecture
 - All main operational requirements and capabilities are available

TIMESTEN SCALEOUT – OUR OPINION

- **TimesTen Scaleout is game-changing**
 - Enables clients to scale up and/or out
 - While getting subsecond response times on all database queries
 - This is very important to our clients!
- **Developers love TimesTen Scaleout!**
 - ACID compliant
 - SQL and PL/SQL compatible
 - Manageable via SQL Developer
- **Deploying and maintaining TimesTen Scaleout is fast and easy**
 - Deployment tools are well documented and easy to use

HOW TO DEPLOY THE TECHNOLOGY SUCCESSFULLY

- **(Early) architectural decisions really matter!**
- **Focus on system and capacity requirements up front**
 - At least the critical requirements should be well known and approved in advance
- **Testing is at least as important as development**
- **Work together with us or Oracle when designing the system**
 - Don't wait until implementation time!
 - Benefit from our partnership and past experience.
 - As well, we will keep in mind features that hopefully come in future as k>3 or disaster recovery!

Q

&

A

Integrated Cloud

Applications & Platform Services