

ORACLE®

ORACLE®
TimesTen

ORACLE®
TimesTen Scaleout

Oracle TimesTen In-Memory Database Overview

Sam Drake

Architect, Oracle TimesTen In-Memory Database

 **In-Memory
Computing**
SUMMIT | NORTH
AMERICA
2018

ORACLE®

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

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.

Best In-Memory Databases: For Both OLTP and Analytics


In-Memory for OLTP

Oracle Database In-Memory

Database

Andy Rivenes 2:40pm TODAY Bayside Room

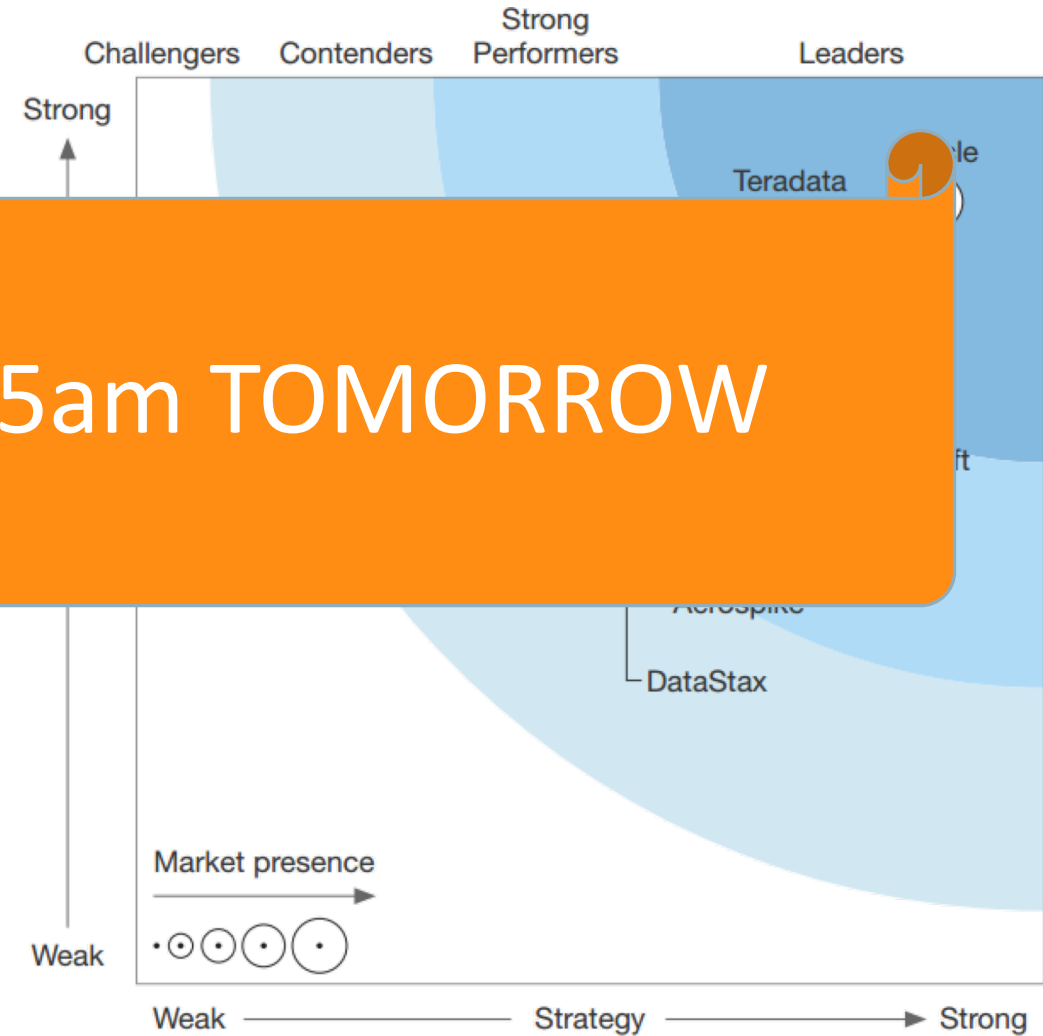
In-Memory for Analytics



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



Tirthankar Lahiri 10:05am TOMORROW

<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.

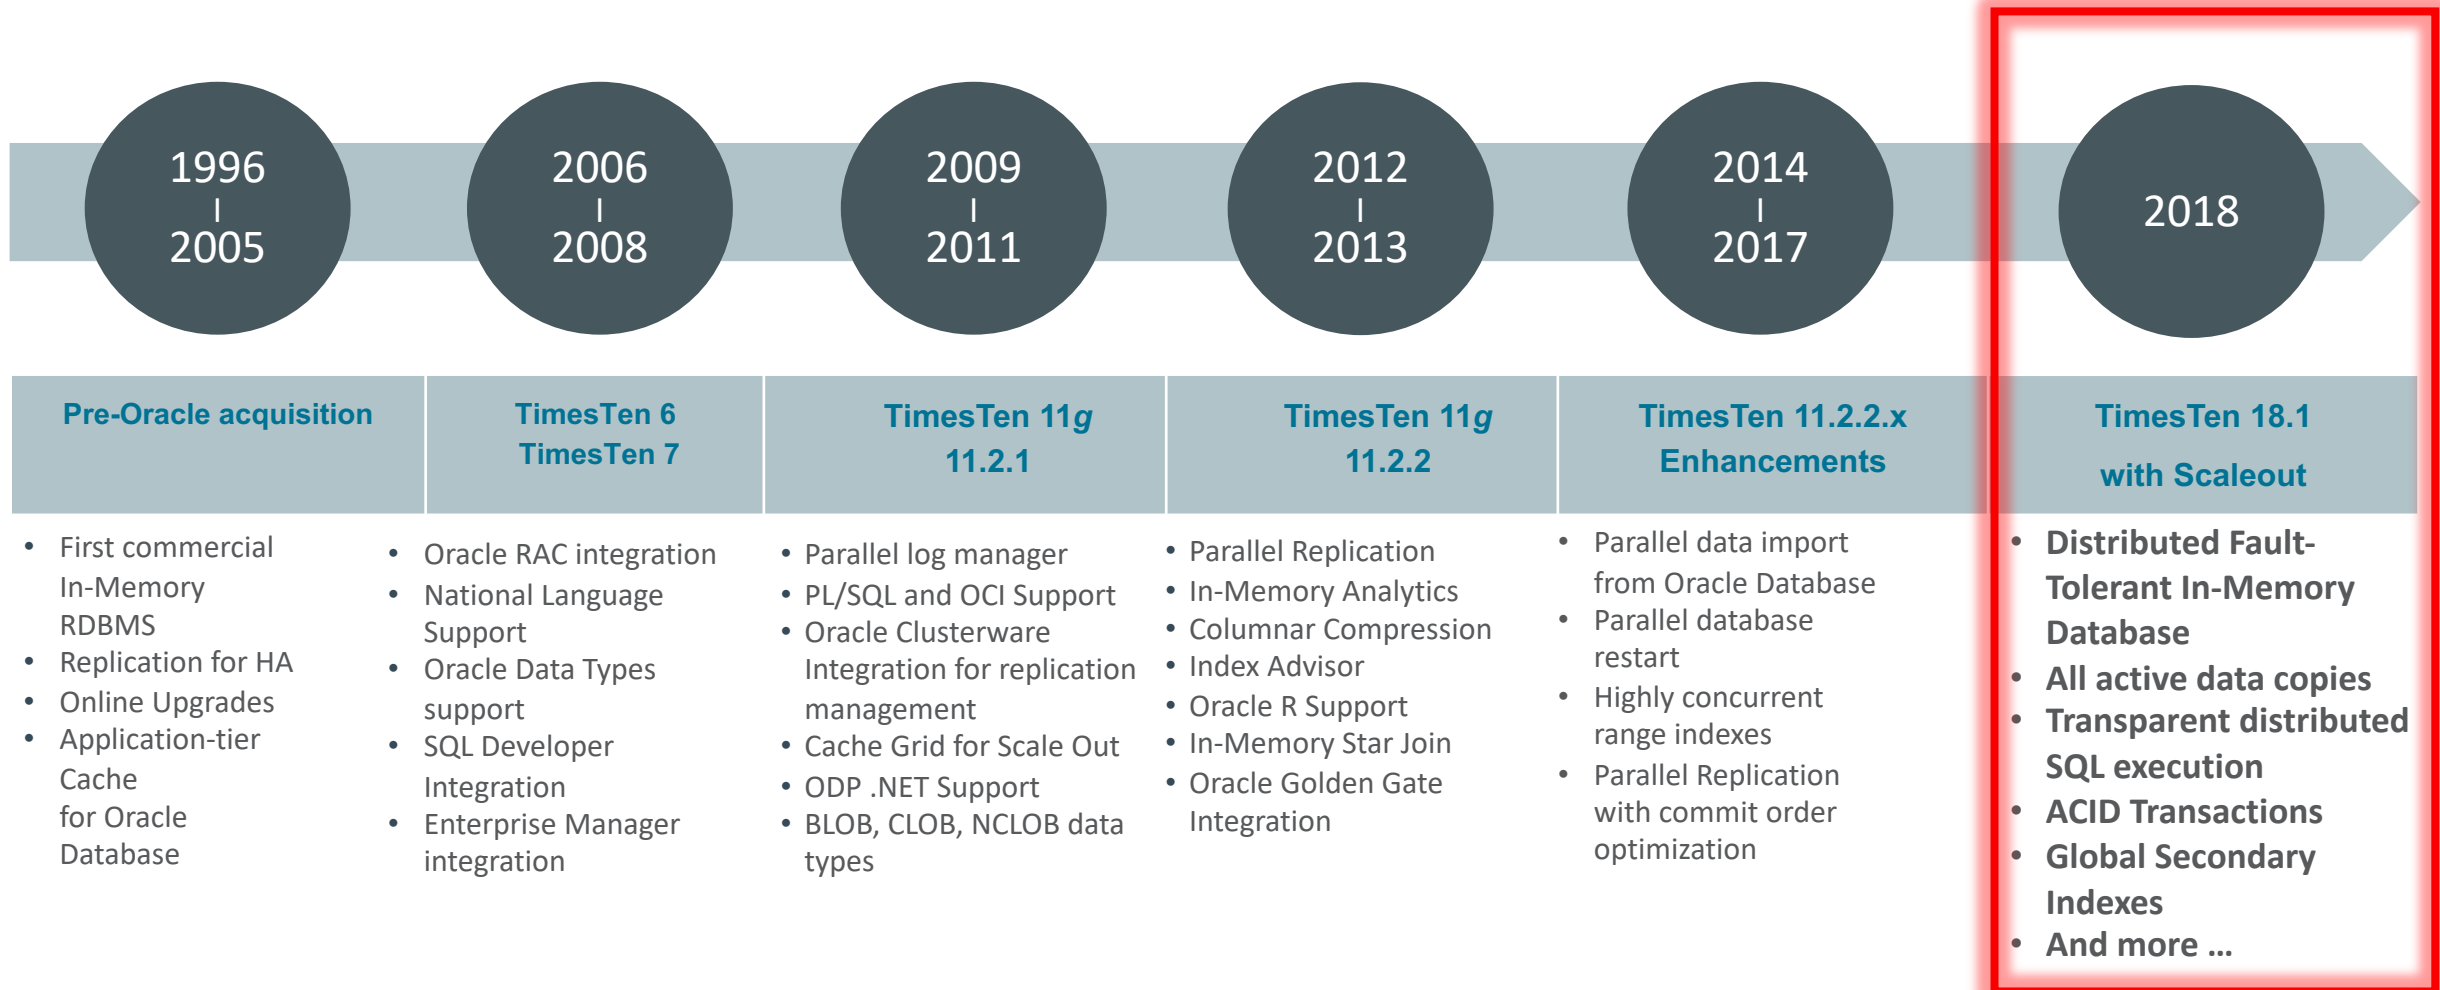
Agenda

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

Introduction

Oracle TimesTen – Class Leading In-Memory Database

20+ Years of Extreme Performance



Most Widely Used Relational In-Memory Database

Deployed by Thousands of Companies



Oracle TimesTen In-Memory Database

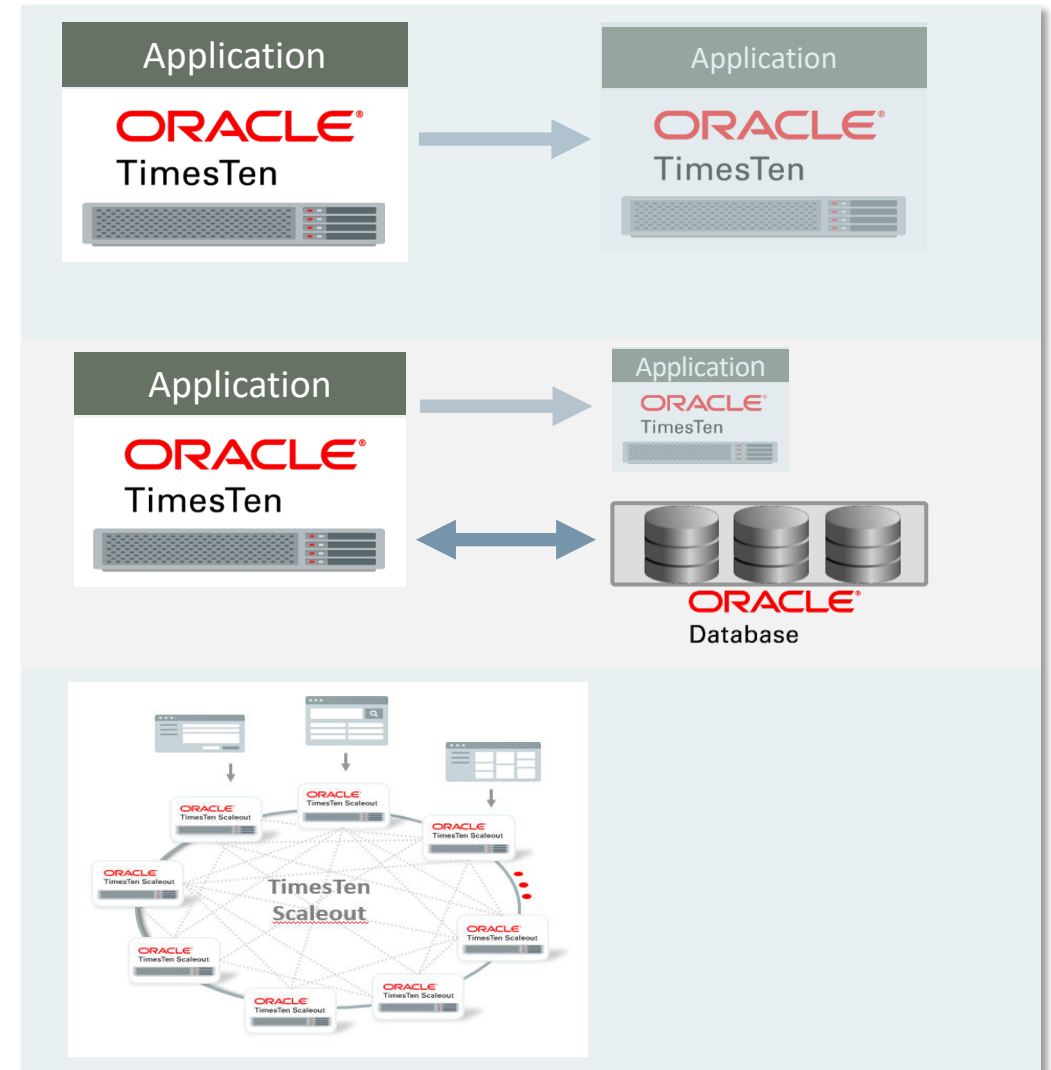
Multiple Deployment Options

TimesTen Classic

- Standalone / Replicated Relational IMDB
- Cache for Oracle Database
Microsecond response time, millions of TPS throughput

TimesTen Scaleout

- Distributed Relational IMDB
Hundreds of millions of TPS throughput



TimesTen Classic


TimesTen Classic

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
- HA and Disaster Recovery

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

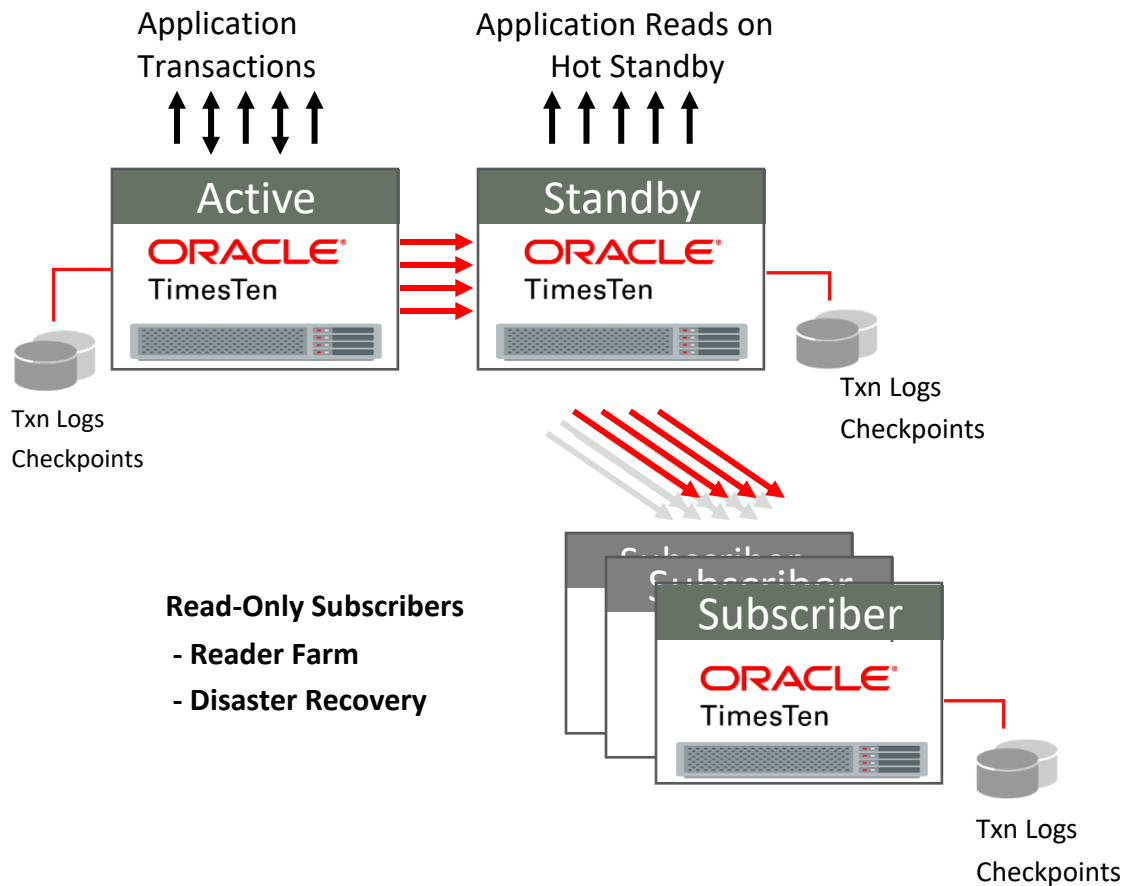
5.6 Million Transactions Per Second

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)



Real-Time Transactional Replication

High Availability and Disaster Recovery

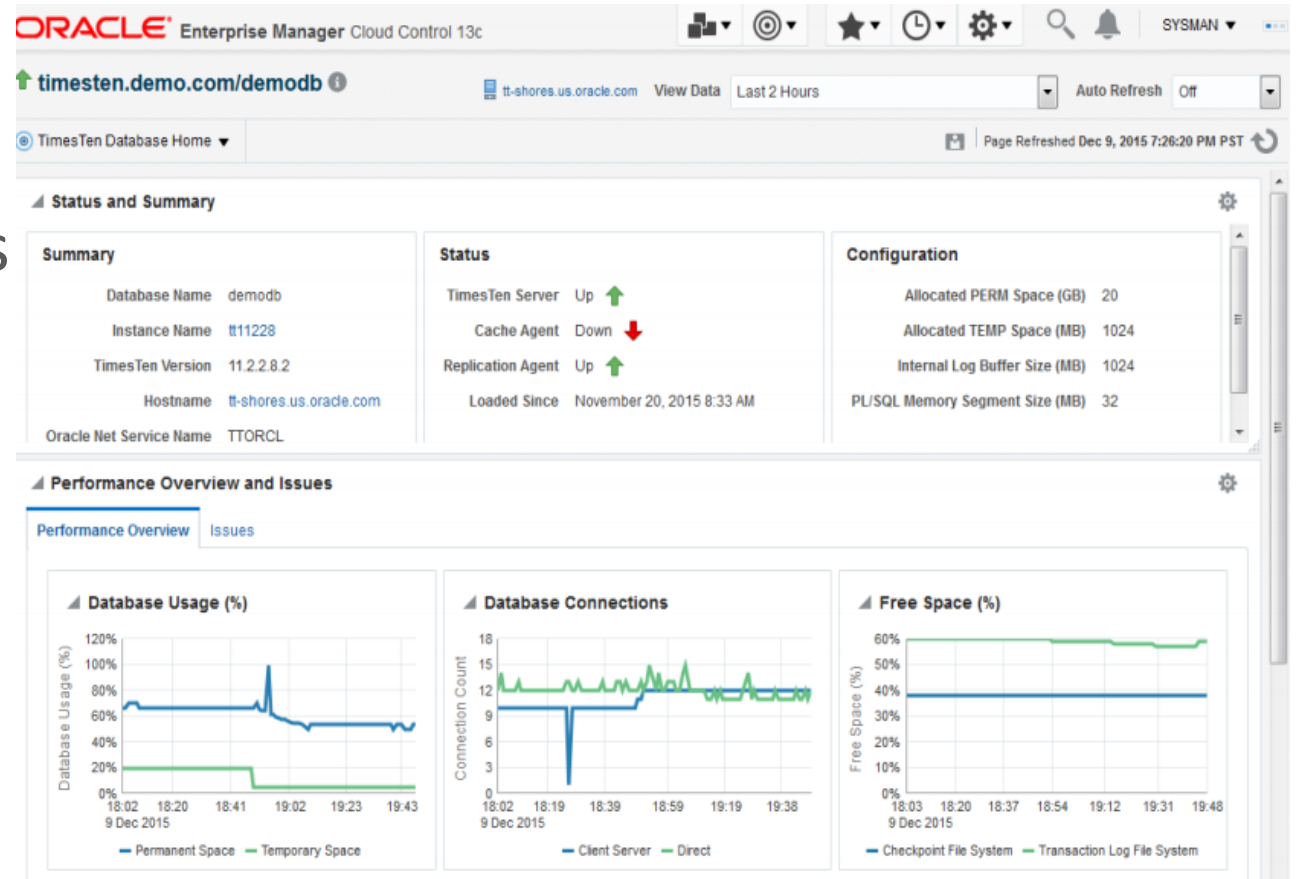


- High performance
 - Synchronous / Asynchronous
 - Parallel send of log streams
 - Parallel apply of changes on Standby and Subscribers
- HA and DR support
- Online rolling upgrades
 - No application downtime
 - Cross-version replication
- Integration with Oracle Clusterware


Enterprise Manager for TimesTen

System Monitoring Plug-in

- Real-time performance and availability monitoring
- TimesTen databases and instances administration
- Automate backups and restore
- Cache and Replication activity reporting
- SQL and Transaction monitoring



Application Development



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

ORACLE®
TimesTen

Coming soon:

The 'Coming soon' section is enclosed in a red dashed border and features logos for several emerging technologies: node.js, python, GO (with a blue character), REST APIs (with a green curly brace icon), Ruby (with a red gem icon), and a red database icon with an upward arrow.

Ericsson Mobile Positioning System (MPS)



Application Overview

- Industry : Telecom
- Business : Business&Operation Support System
- Application : Mobile Positioning System
 - GMPC node of MPS collects and utilizes mobile subscribers' location information
 - MPS has 120+ installations distributed over Americas, Europe, Asia-Pacific and Africa

Challenges

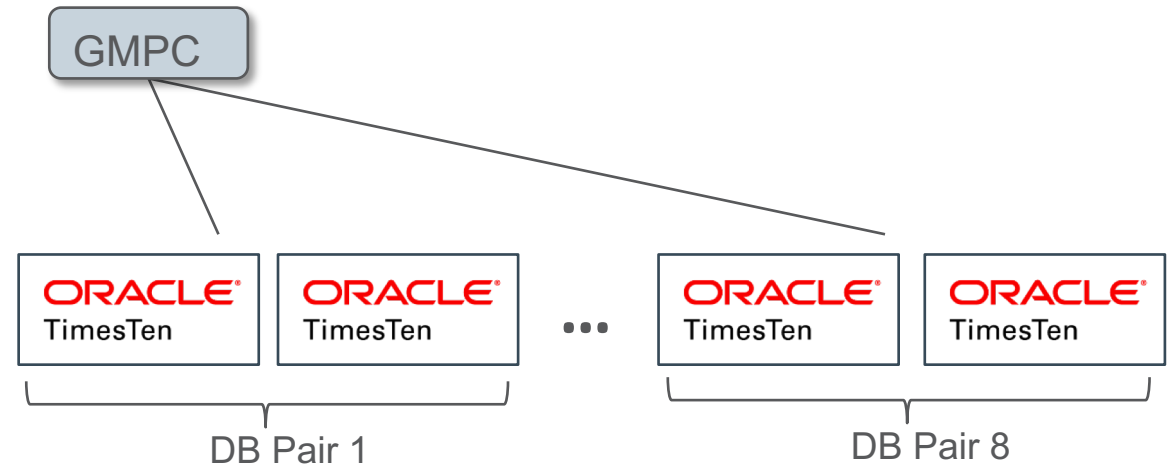
- Need highly concurrent mobile locations updated randomly
- Need high transaction throughput with consistent low latency
- Need 24x365 availability

Solution

- Oracle TimesTen Database
- Oracle TimesTen Replication for High Availability

Why TimesTen ?

- End-to-end response time ~1.5 milliseconds
- 63,000 transactions per second (replicated in real time) per database
- Multi-DB sharded architecture to achieve increased performance – about 252k transactions per second
 - Partition the data across multiple TimesTen active-standby database pairs
 - The number of database server pairs is 8



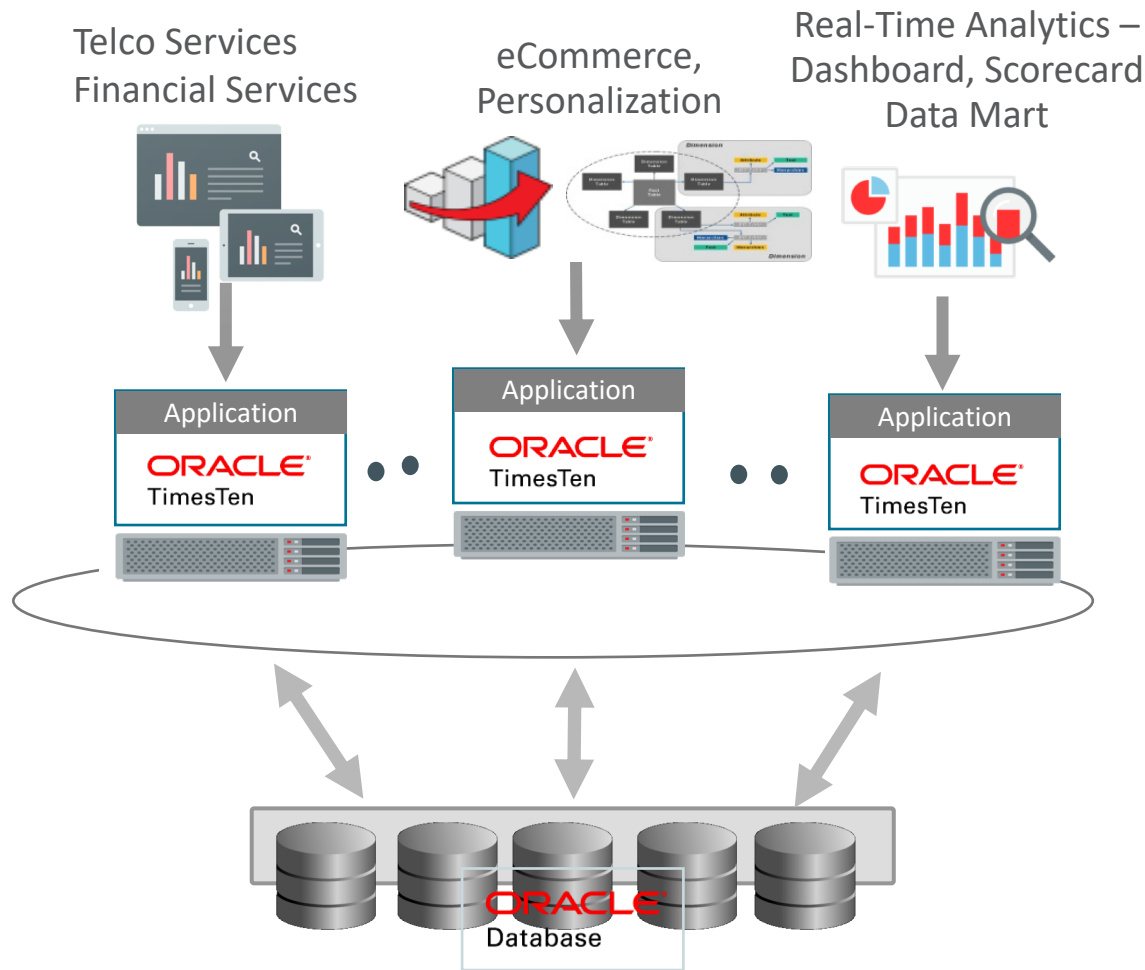
TimesTen Classic Summary

- 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 with standard APIs
- Compatible with Oracle products
Oracle Enterprise Manager, SQL Developer, Oracle GoldenGate, Oracle Clusterware

TimesTen Application-Tier Database Cache


For Oracle Database

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

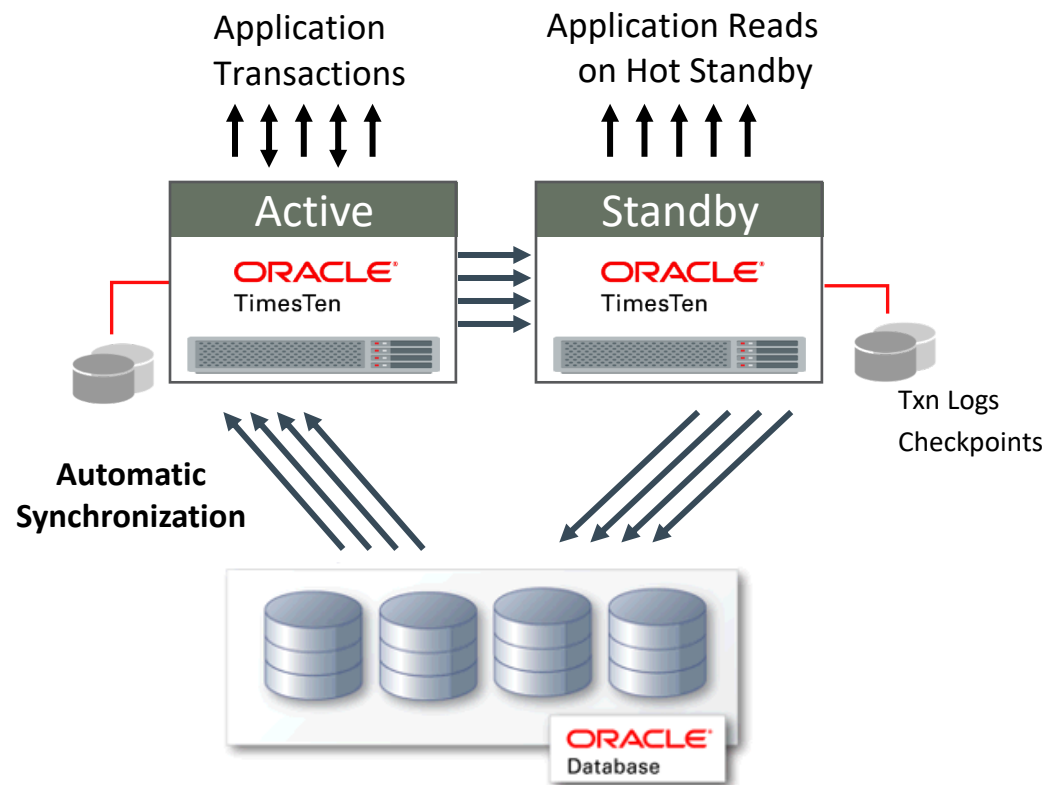
Flexible Cache Group Configurations



- 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



- 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

# Phone Agent Task Assignments

## Application Overview

- Industry : Finance
- Business : Insurance, Banking, Investment
- Application : Agent Task Management
  - Automatic task assignment based on predefined rules
  - Manually reassign task from one agent to another

## Challenges


- Database scalability with extreme high concurrency affecting end to end response time
- Maintain user satisfaction
- Minimal changes to existing architecture and application
- Must be highly available

## Solution

- Oracle TimesTen Application-Tier Database Cache
- TimesTen Replication for High Availability
- Oracle Database

## Why TimesTen ?

- Delivered lower and consistent response time; achieved 40x improvement in both response time and throughput
- Automatic data synchronization between TimesTen and Oracle Database
- With built-in HA, supports automatic failover and switchover





# TimesTen Cache Summary

- Accelerating existing Oracle Database Applications

*Caching from Oracle Database and automatic synchronization with Oracle Database*

- Shares the same architecture as TimesTen Classic

*Fast and consistent response time with low latency*

*High availability and online upgrades*

*Standard SQL/relational model with standard APIs*

*Compatible with Oracle Enterprise Manager, SQL Developer, GoldenGate and Clusterware*

- Multiple configuration options – mix and match

*Read-only cache groups*

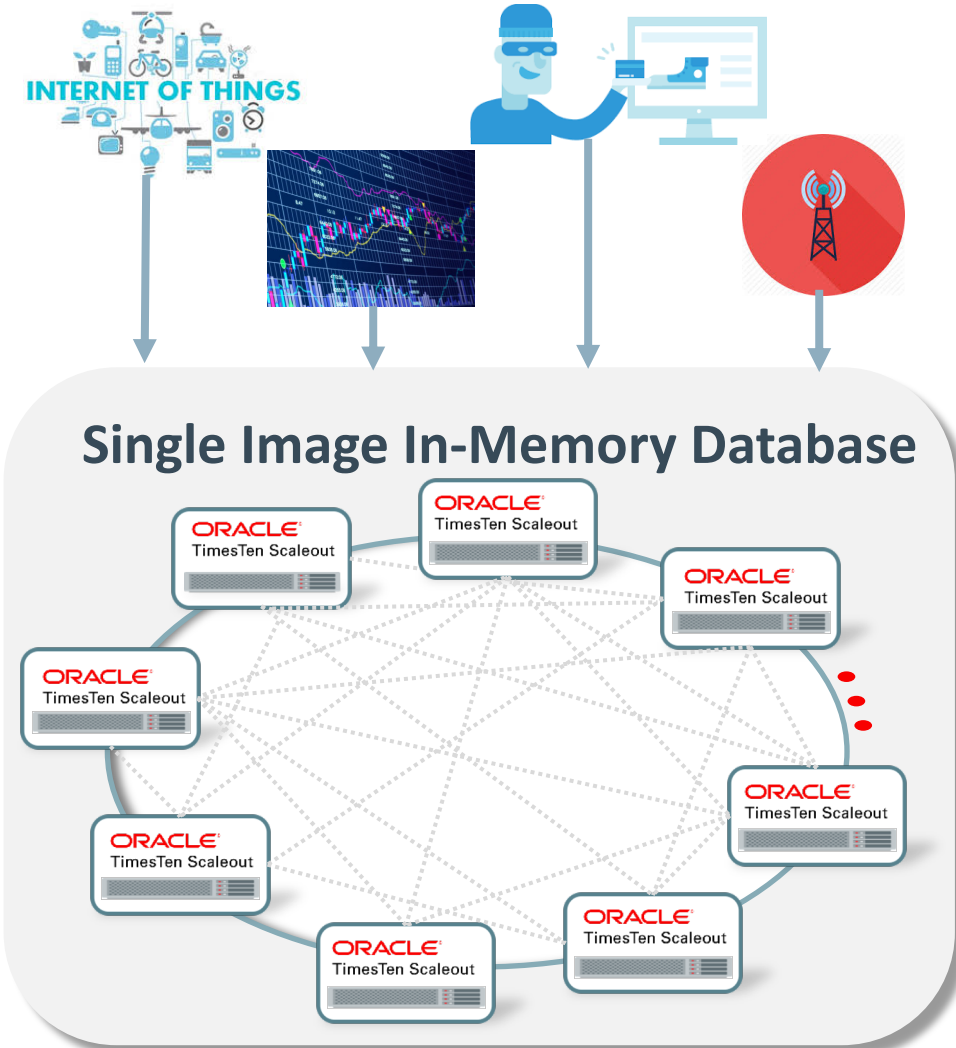
*Writethrough cache groups*

*Native TimesTen tables*

# TimesTen Scaleout

**Distributed, Elastically Scalable, Single Image, Fault Tolerant**

# TimesTen Scaleout

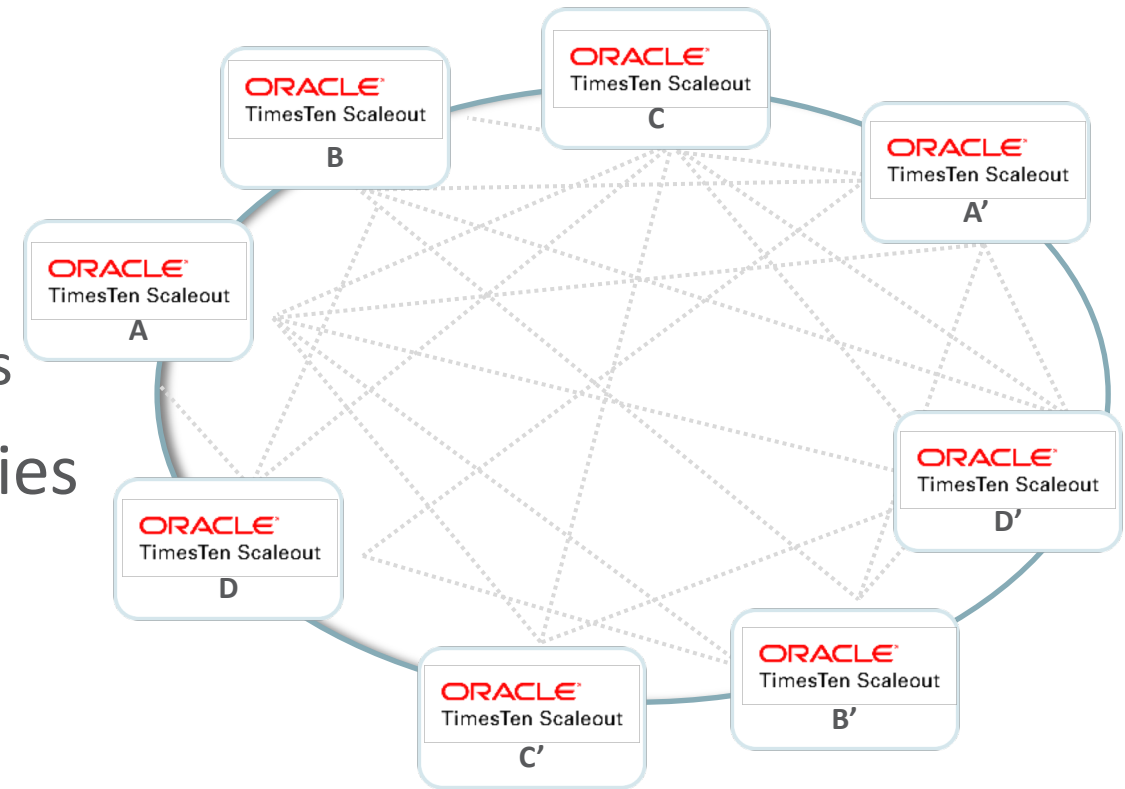


- For High-Velocity **Extreme OLTP** applications
  - IOT, trading, fraud detection, mobile, click stream, billing, orders, etc.
- Cutting-Edge Design:
  - Pure In-Memory, Full SQL, Full ACID Transactions
    - Scale-out shared nothing architecture
    - Multiple data copies for HA (K-safety)
      - All copies active for read/writes
    - Global secondary indexes
  - Complex SQL and Parallel SQL for reporting and batch
- Centralized management and administration

# Distributed, Shared Nothing, In-Memory Database

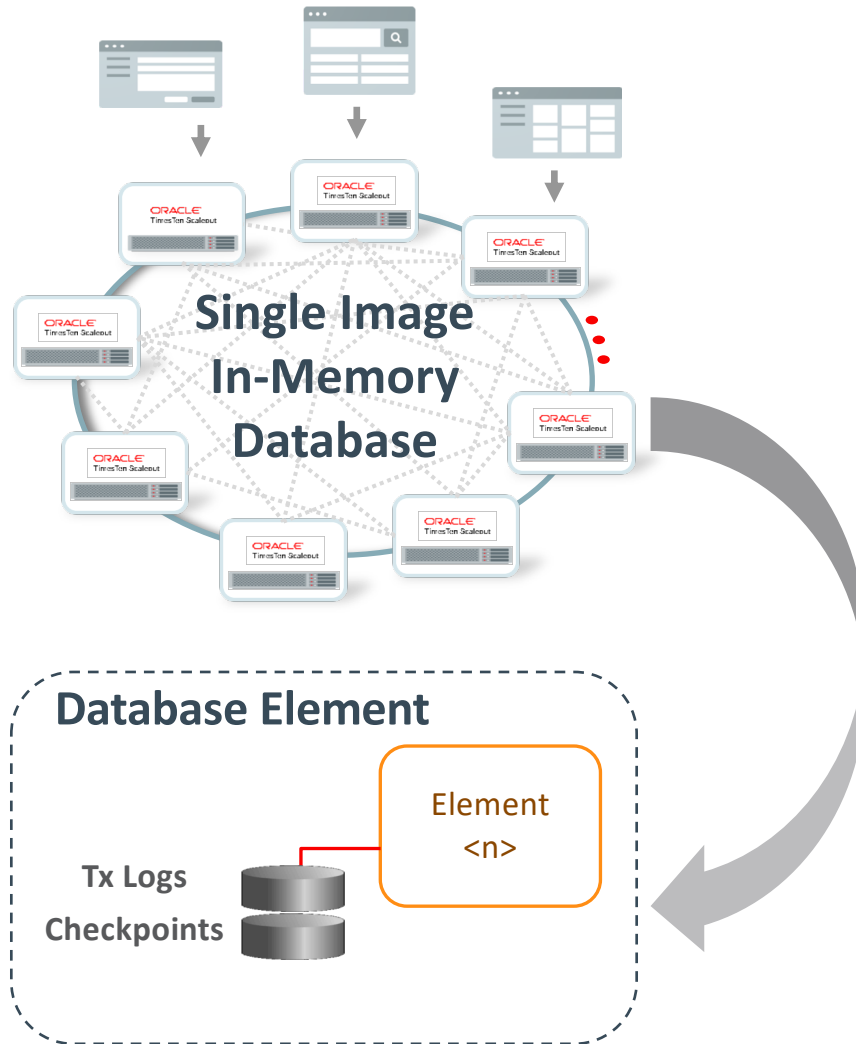
## Single-Image DB with High Availability and Elasticity

- 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



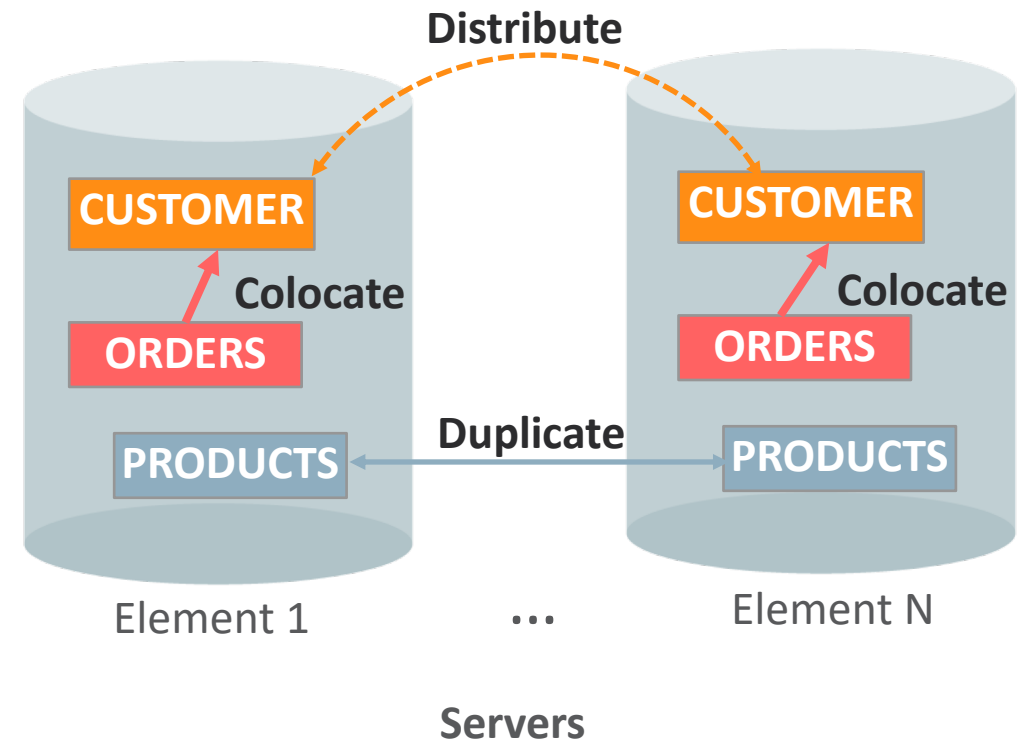
Element 1



Element 2

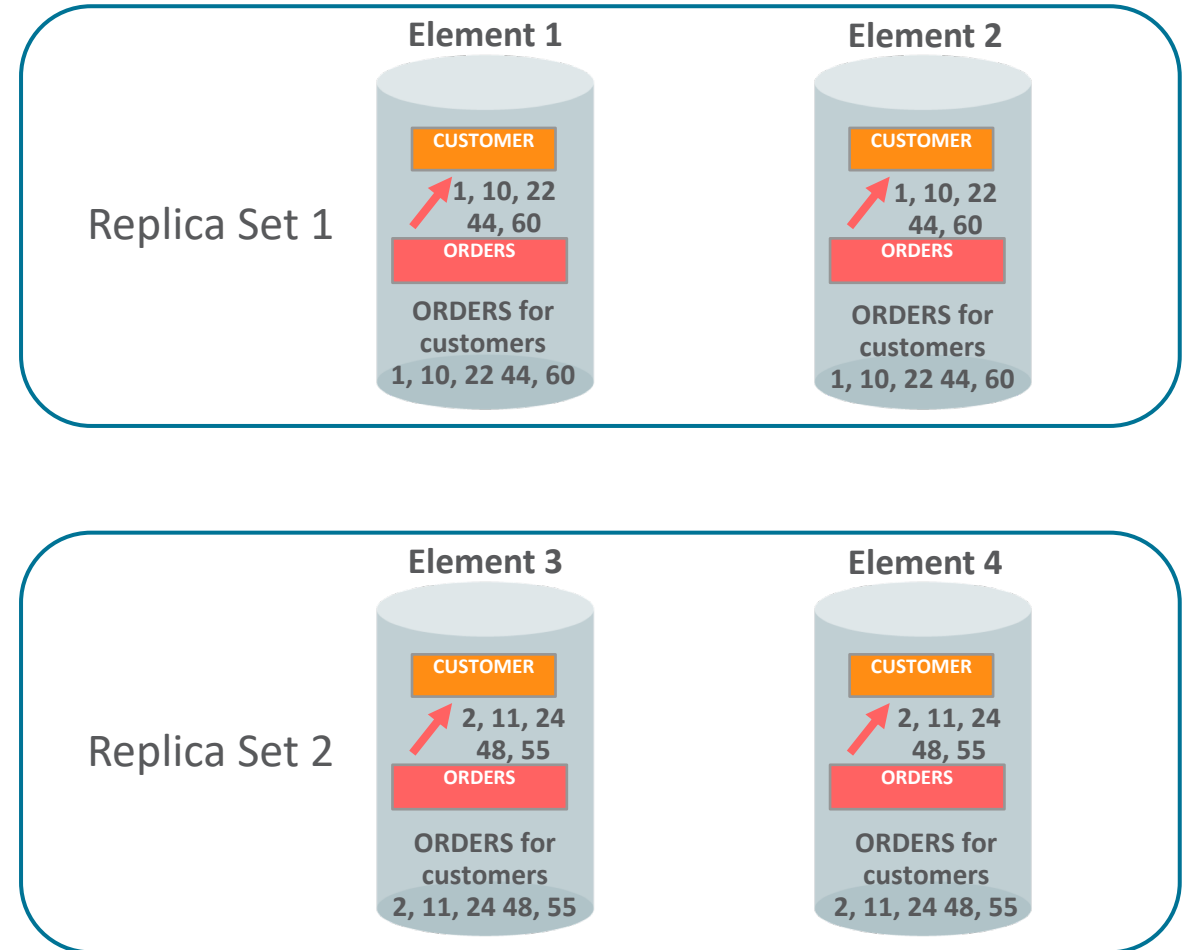
# TimesTen Scaleout - Data Distribution

- **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



# 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
- Both elements are “active”
- Queries and transactions can span any/all replica sets
- Two phase commit protocols keep them in sync



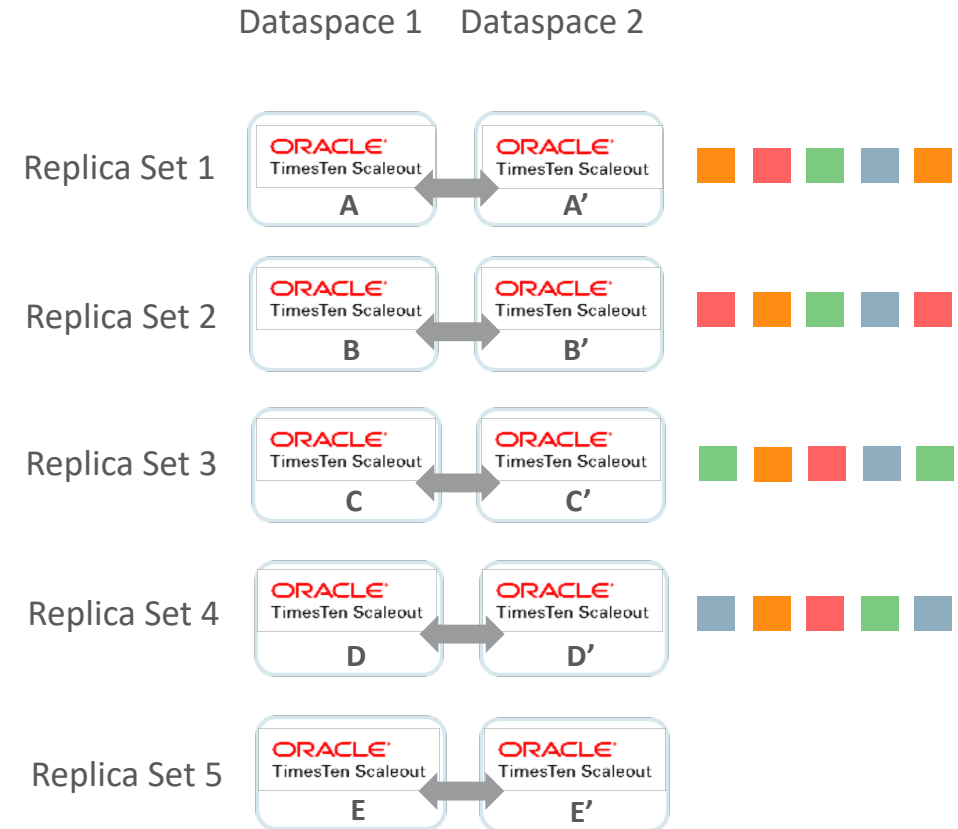


# TimesTen Scaleout - Elastic Scalability

Expand and shrink the database based on business needs

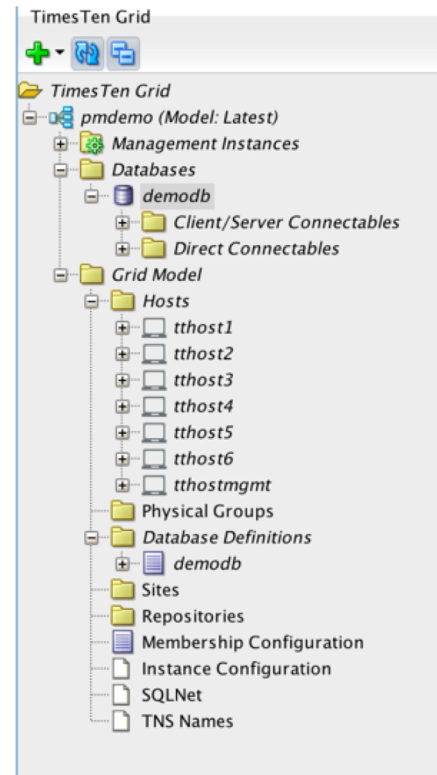
## Adding and removing DB elements

- Data redistributed to new elements
- Workload automatically uses the new elements
- Connections will start to use new elements
- Throughput increases with added 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



TimesTen Grid

TimesTen Grid

pmdemo (Model: Latest)

Management Instances

Databases

demodb

Client/Server Connectables

Direct Connectables

Grid Model

Hosts

tthost1

tthost2

tthost3

tthost4

tthost5

tthostmgmt

Physical Groups

Database Definitions

demodb

Sites

Repositories

Membership Configuration

Instance Configuration

SQLNet

TNS Names

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                |






# TimesTen Scaleout Performance

# DB Comparison : YCSB Workload B (95% read, 5% write)

| Database                                                                                        | TPS   | Nodes |
|-------------------------------------------------------------------------------------------------|-------|-------|
|  mongoDB       | 260 K | 2     |
|  redis         | 1 M   | 3     |
|  apache ignite | 454 K | 4     |
| <b>VOLTD</b> DB                                                                                 | 1.5 M | 6     |
|  Couchbase    | 454 K | 9     |
|  cassandra   | 640 K | 30    |

**ORACLE**<sup>®</sup>  
TimesTen Scaleout ?

# DB Comparison : YCSB Workload B (95% read, 5% write)

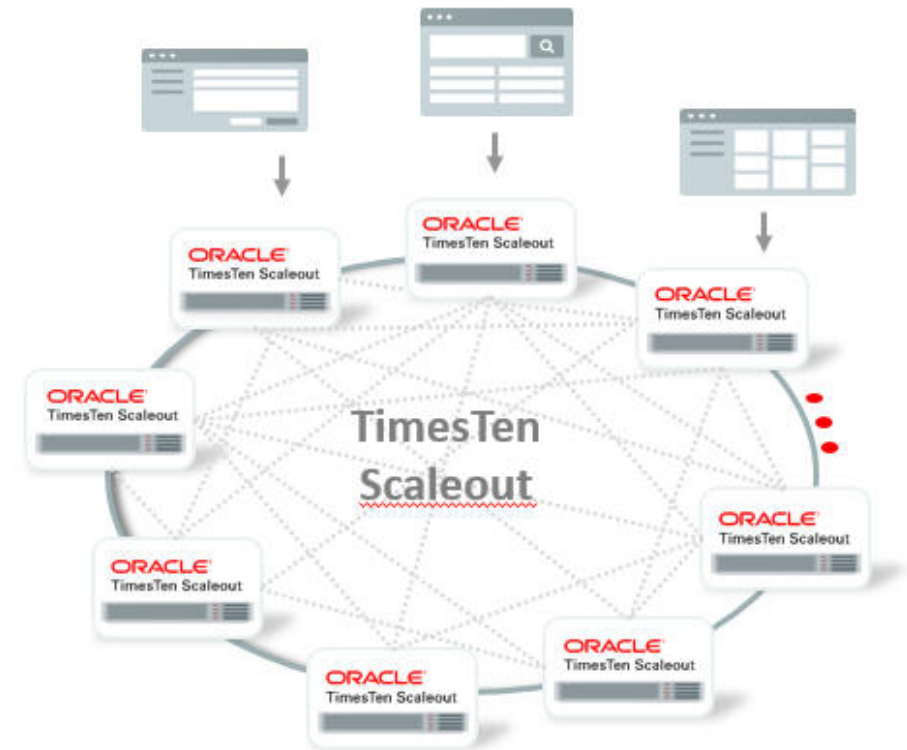
| Database                                                                                        | TPS   | Nodes |
|-------------------------------------------------------------------------------------------------|-------|-------|
|  mongoDB       | 260 K | 2     |
|  redis         | 1 M   | 3     |
|  apache ignite | 454 K | 4     |
| <b>VOLTD</b> DB                                                                                 | 1.5 M | 6     |
|  Couchbase    | 454 K | 9     |
|  cassandra   | 640 K | 30    |

| Database                           | TPS           | Nodes     |
|------------------------------------|---------------|-----------|
| <b>ORACLE</b><br>TimesTen Scaleout | <b>2.7 M</b>  | <b>2</b>  |
| <b>ORACLE</b><br>TimesTen Scaleout | <b>5.5 M</b>  | <b>4</b>  |
| <b>ORACLE</b><br>TimesTen Scaleout | <b>10.6 M</b> | <b>8</b>  |
| <b>ORACLE</b><br>TimesTen Scaleout | <b>38 M</b>   | <b>32</b> |



# TimesTen Scaleout Summary

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



Q

&

A

# Integrated Cloud

## Applications & Platform Services