


ORACLE®

Oracle TimesTen In-Memory Database 18.1 Scaleout Functionality, Architecture and Performance

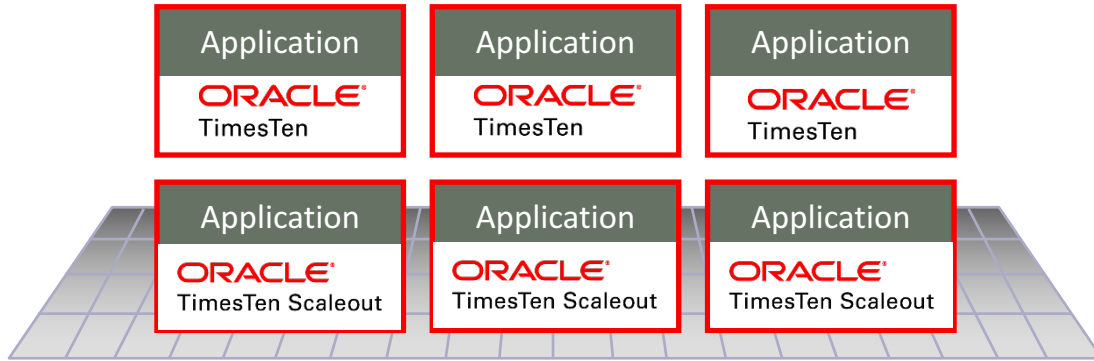
Chris Jenkins

Senior Director, In-Memory Technology
TimesTen Product Management




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

Oracle TimesTen – Class Leading In-Memory Database

20+ Years of Extreme Performance and Reliability



Most Widely Used Commercial Relational IMDB

Deployed by thousands of companies



Oracle TimesTen In-Memory Database

Multiple Deployment Options

TimesTen Classic

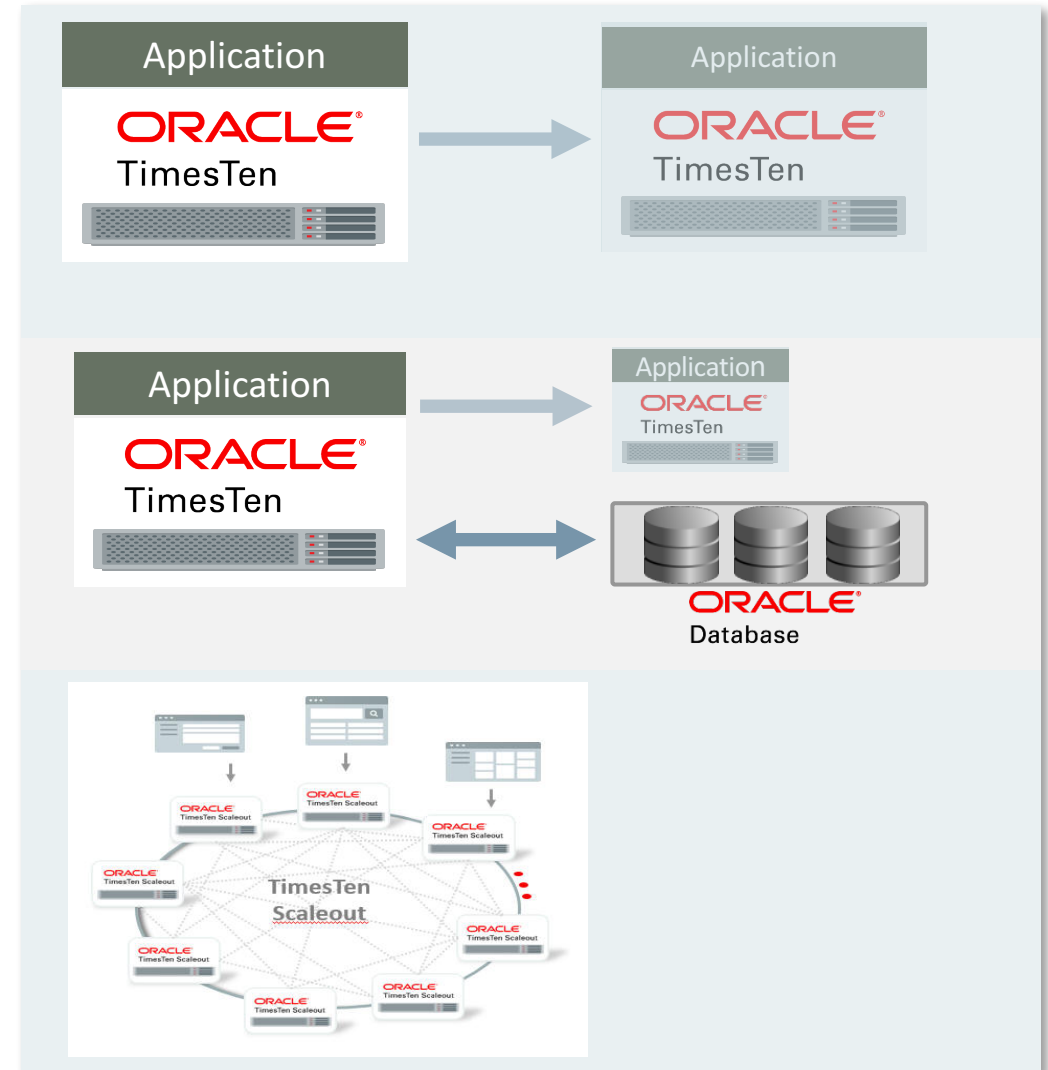
1. Standalone / Replicated Relational IMDB
 - Low latency/high throughput applications
 - ISV/OEM embedded solutions
2. Cache for Oracle Database
 - Accelerate Oracle Database OLTP applications
 - HA option via TimesTen Replication

Microsecond response time, millions of TPS throughput

TimesTen Scaleout – *new in 18.1*


3. Distributed, Shared Nothing, Relational IMDB
 - High throughput and storage capacity
 - Transparent data distribution
 - Elastic scalability
 - Fault tolerant

Hundreds of millions of TPS throughput




TimesTen Scaleout Benefits

Fast & Scalable




- Low latency and high throughput
- Add elements online to increase throughput and capacity

Fault Tolerant



- No single point of failure
- Fully persistent
- K-Safety for HA
- All elements active for reads and writes

Single DB image



- Appears as a single database to applications
- Data location transparency
- Connect to any element and access all data
- No need to de-normalize

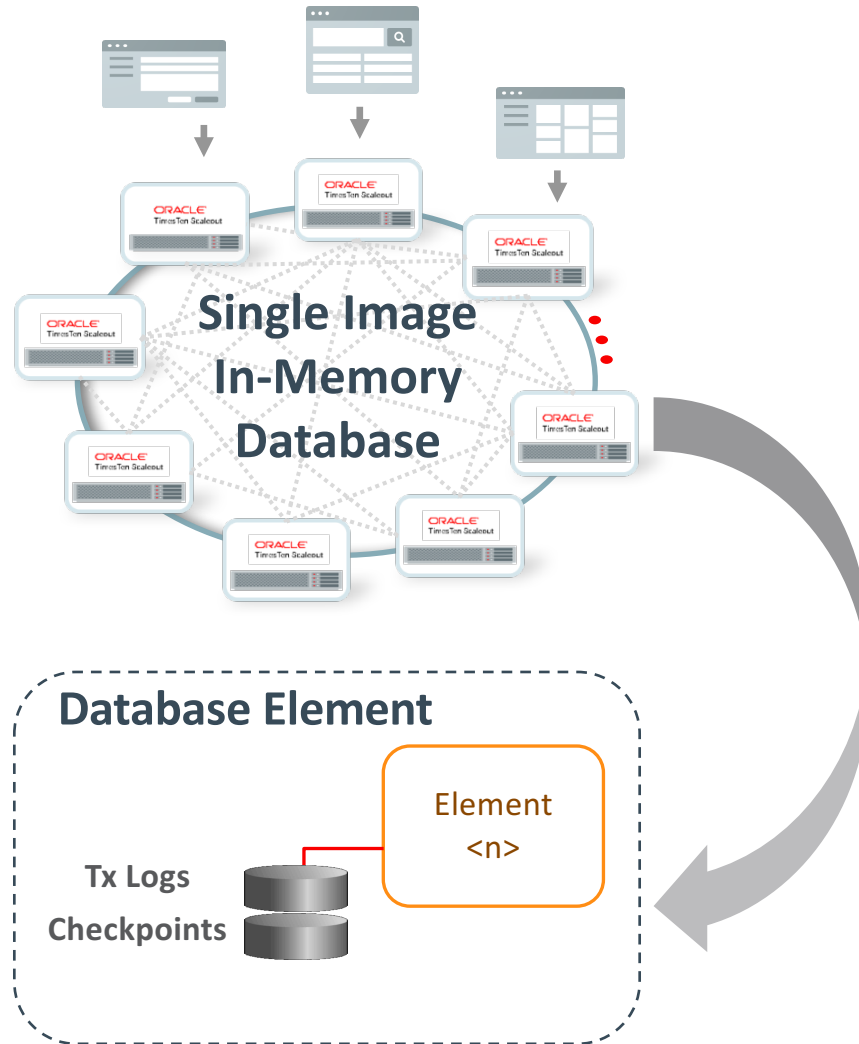
Easy to develop for, Easy to deploy



- Highly compatible with Oracle data types, SQL & PLSQL
- Many languages and APIs
- Monitoring and management from a single location
- Deploy on premises or in the Cloud

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 & recovery
- 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

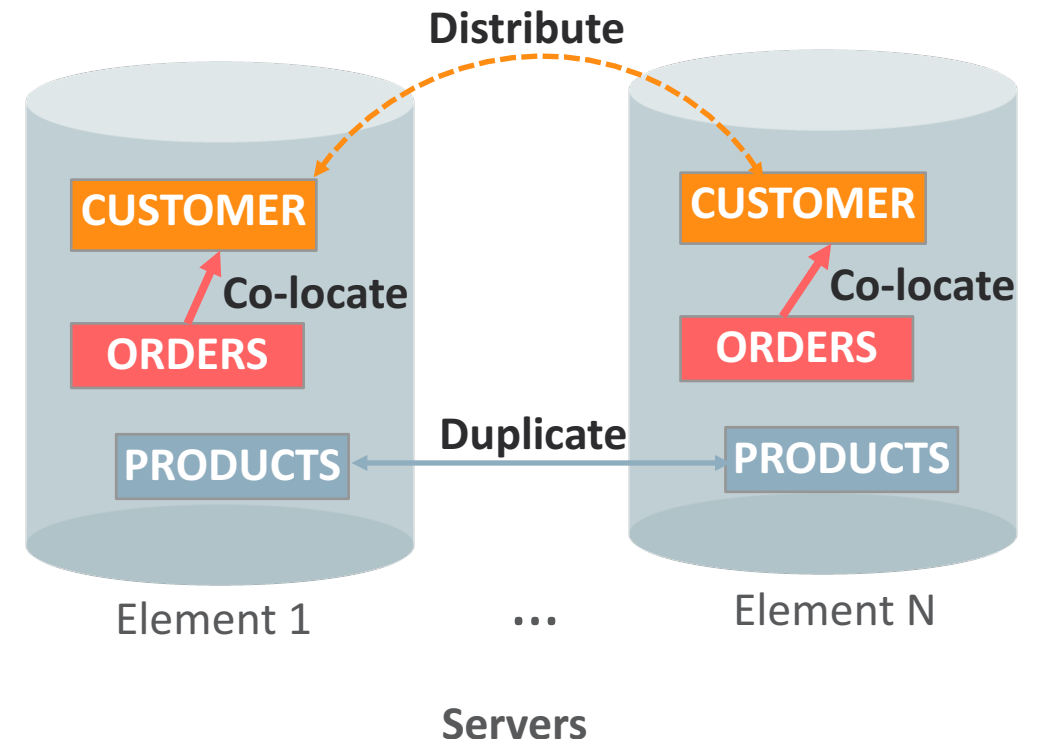


Element 2

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 increase locality of reference
 - 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

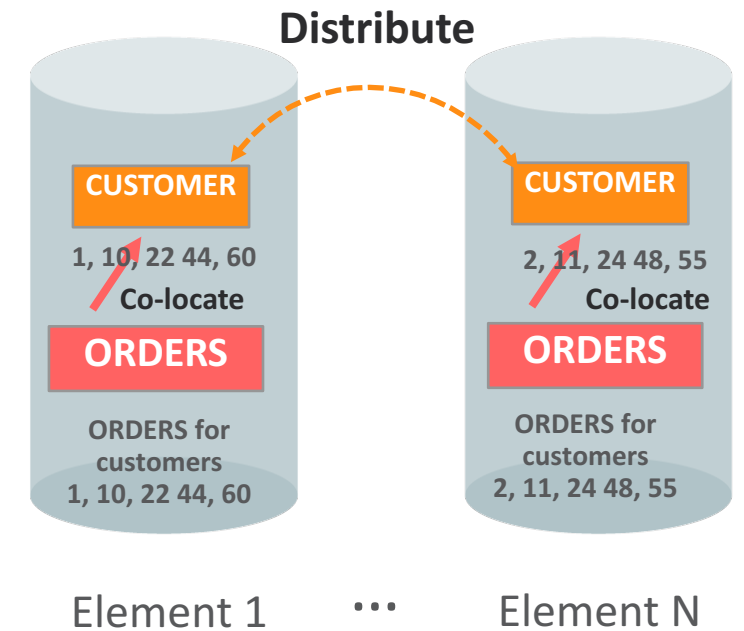
```
CREATE TABLE CUSTOMER (  
  ID NUMBER NOT NULL PRIMARY KEY,  
  NAME VARCHAR2(100)  
) DISTRIBUTE BY HASH;
```



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
- Multiple levels supported

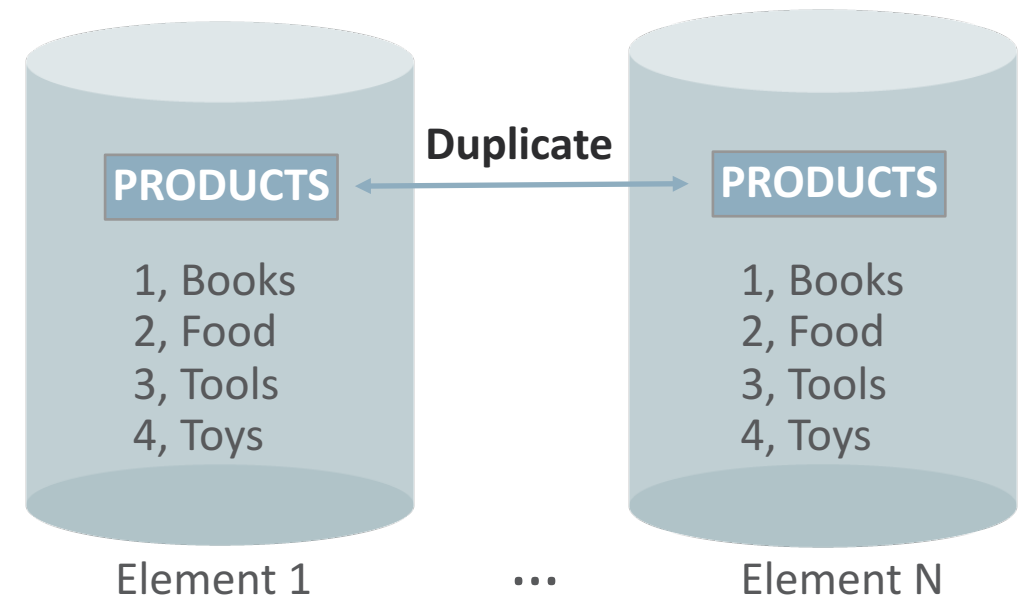
```
CREATE TABLE CUSTOMER (  
  ID NUMBER NOT NULL  
    PRIMARY KEY,  
  NAME VARCHAR2(100)  
  ...  
) DISTRIBUTE BY HASH;  
  
CREATE TABLE ORDERS (  
  ID NUMBER NOT NULL  
    PRIMARY KEY,  
  CUST_ID NUMBER NOT NULL,  
  ...  
  CONSTRAINT FK_CUST  
    FOREIGN KEY (CUST_ID)  
    REFERENCES CUSTOMER(ID),  
  ...  
) DISTRIBUTE BY  
  REFERENCE (FK_CUST);
```



Duplicate

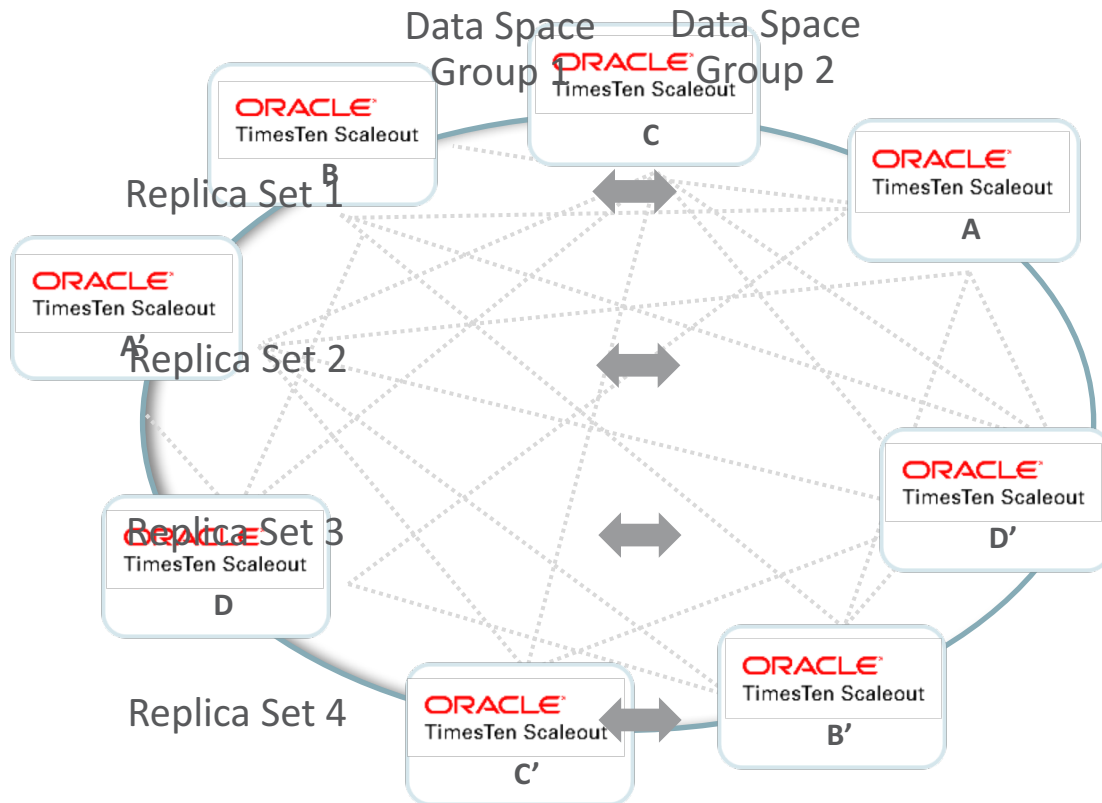
- Every row is present in every element of the grid
- Appropriate for tables that are:
 - Relatively small
 - Frequently read
 - Infrequently modified

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



TimesTen Scaleout - High Availability

K-safety, All Active

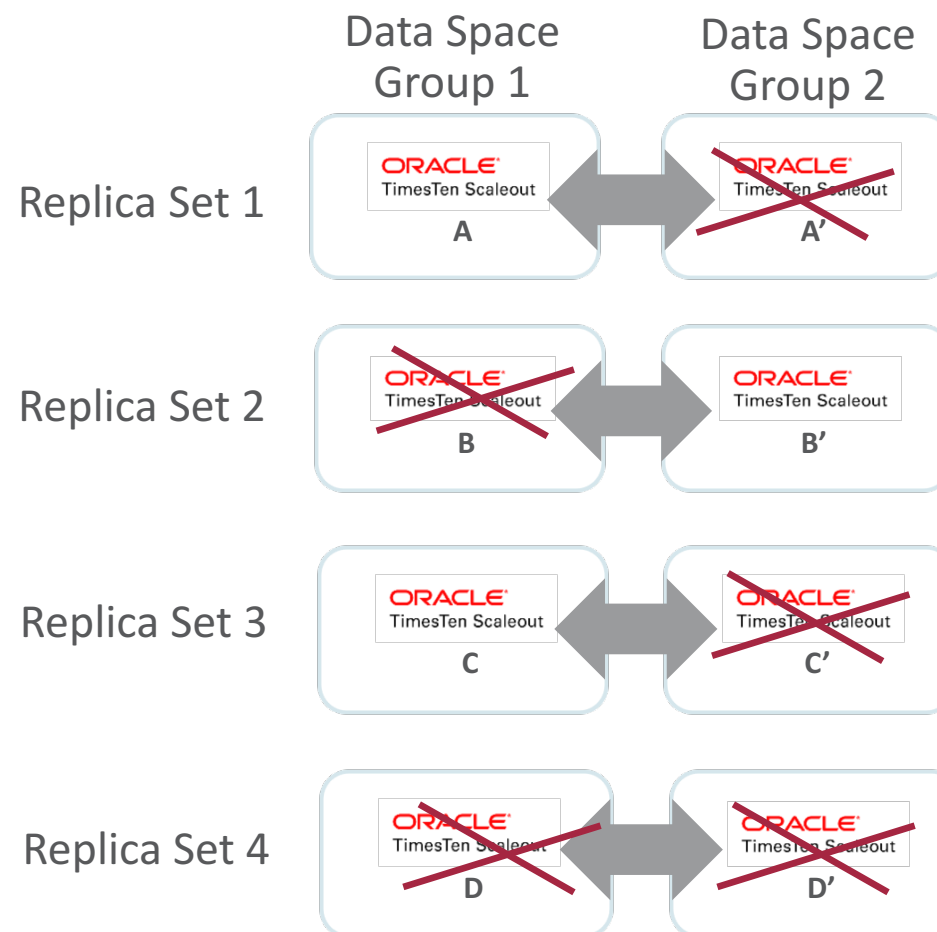


- Built-in HA via multiple copies of the data (K-safety)
 - Automatically kept in sync
- **All** replicas are **active** for **reads** and **writes**
 - Increases the compute capacity
- Transactions can be initiated from, and executed on, any replica
- Queries and transactions can span any/all elements

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
 - Automatic failover for C/S connections
- 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

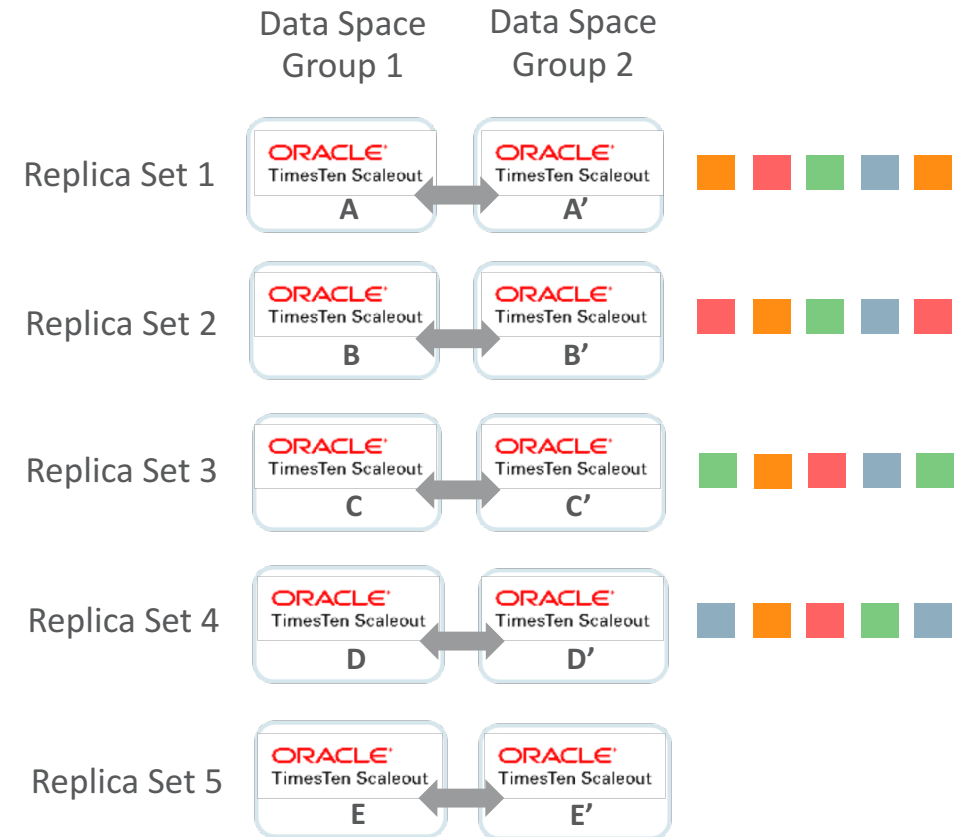


TimesTen Scaleout - Elastic Scalability

Expand and shrink the database based on business needs

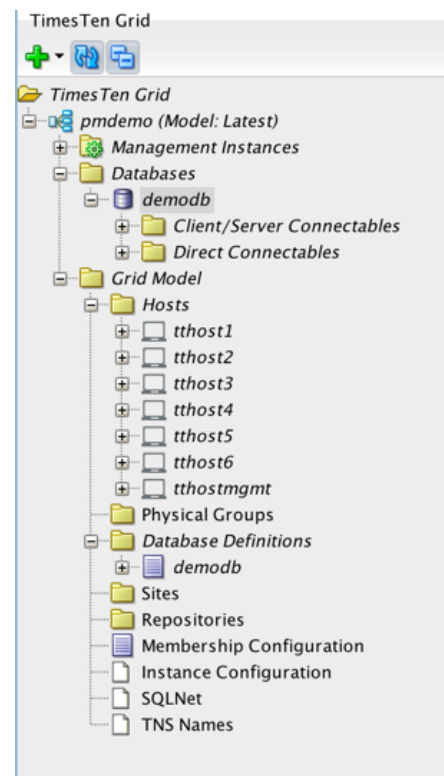
Adding (and removing) replica sets

- 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 are performed from a single host
 - Installing software
 - Patching software
 - Configuration
 - Database creation and management
 - Backup and restore
 - Monitoring
 - Collecting diagnostics
- Command line interface
- SQL Developer (GUI) interface



Status

Database Definition

Topology

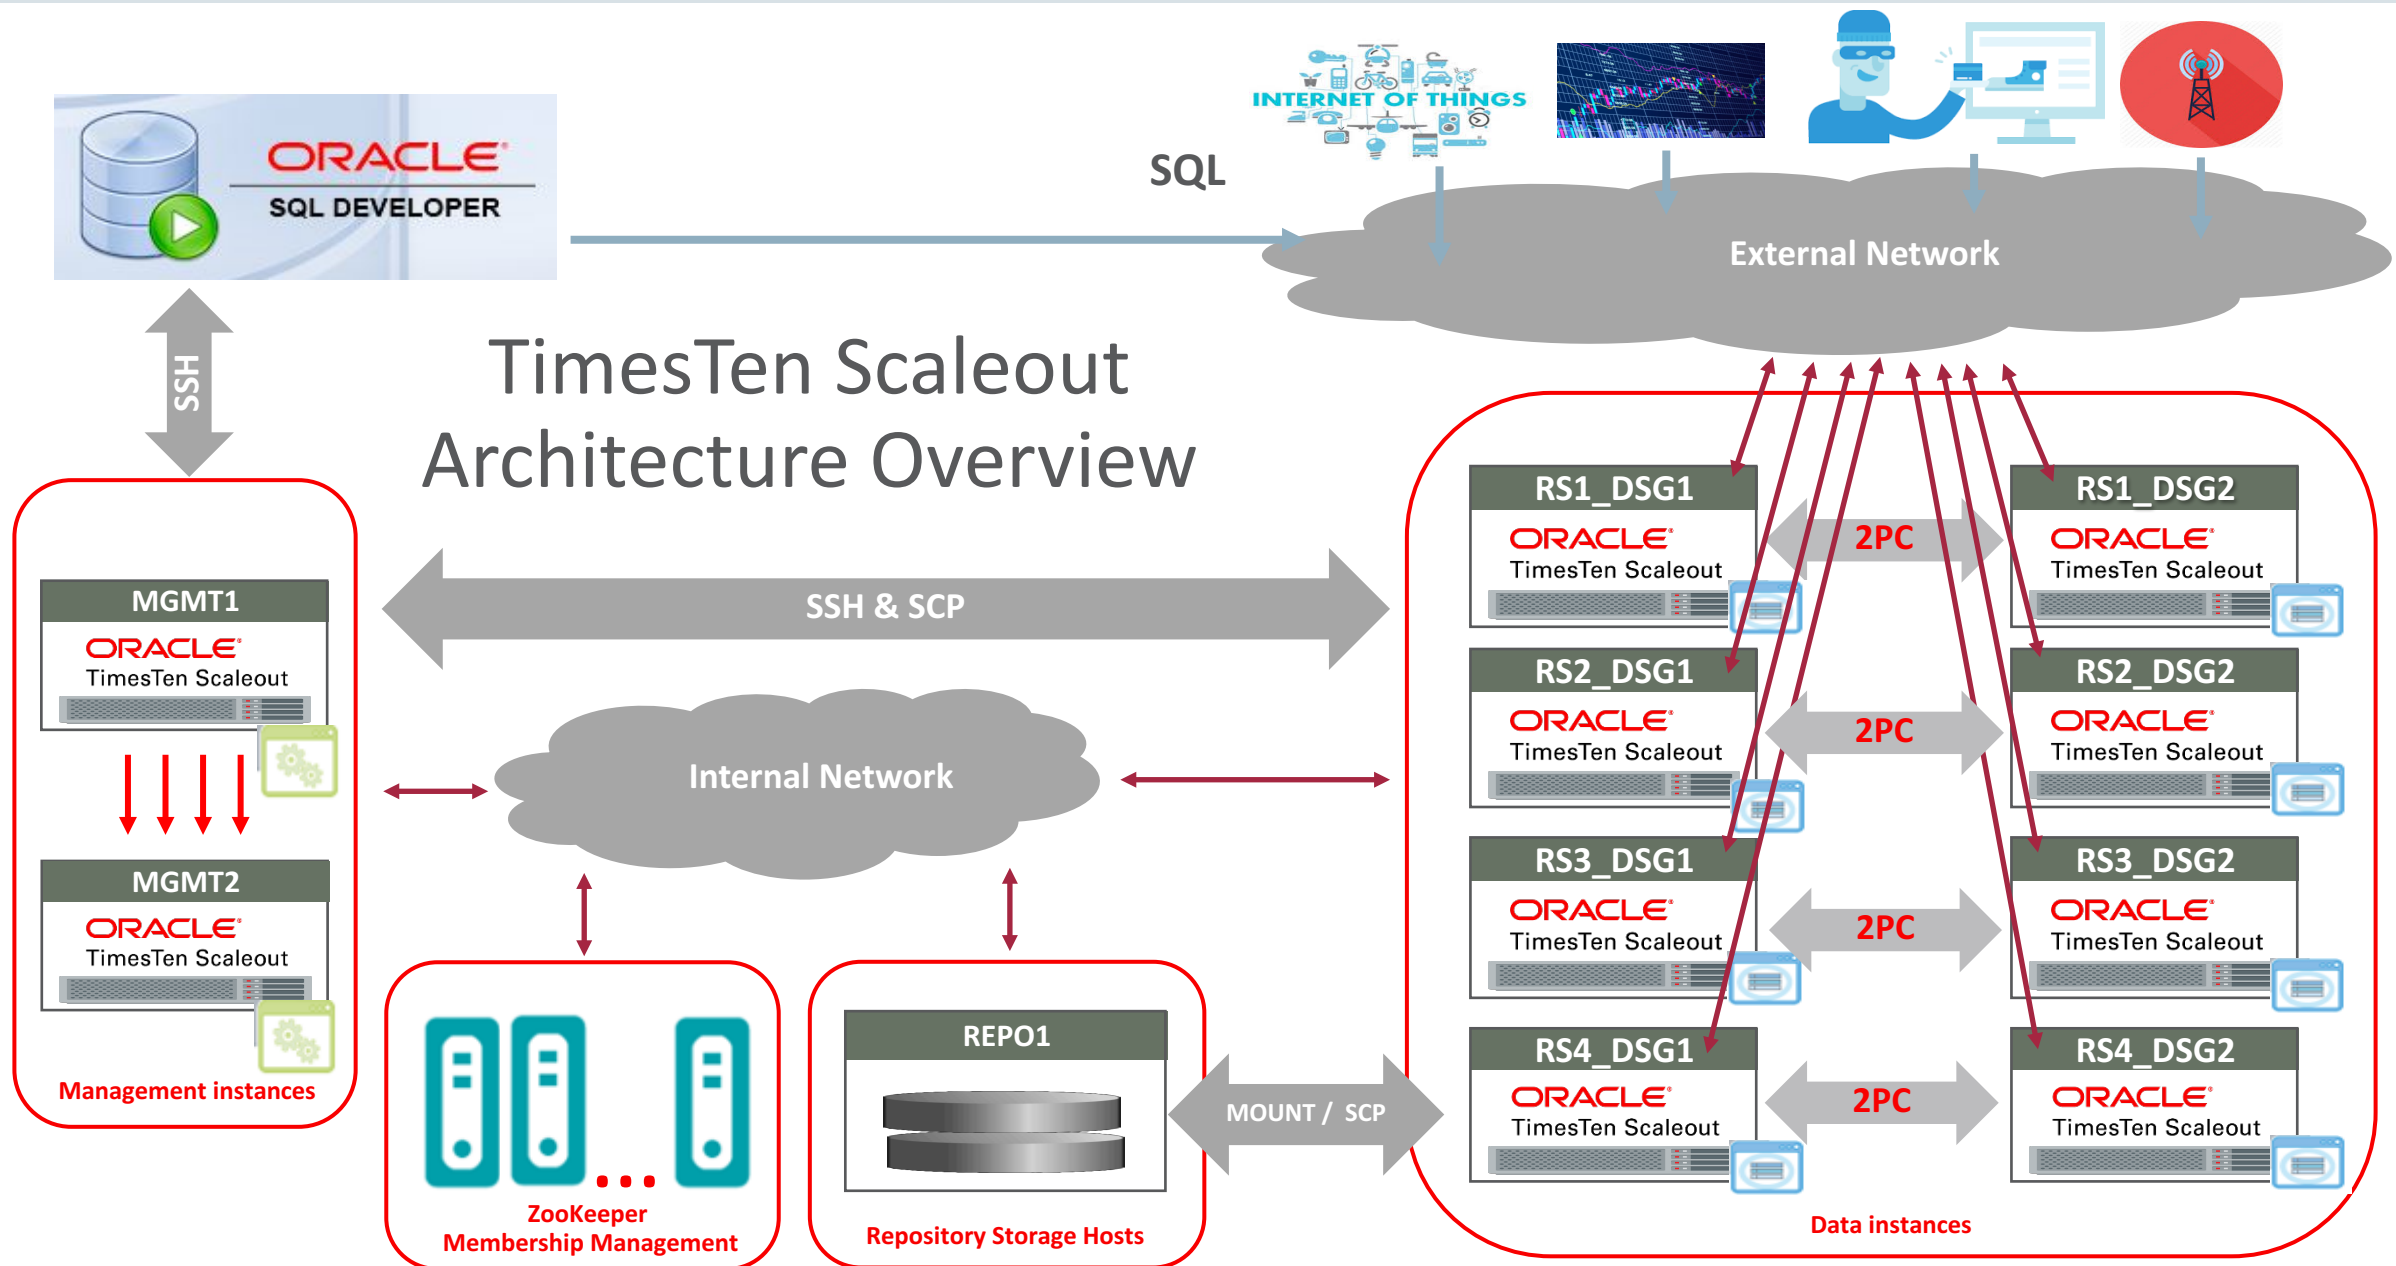
Database demodb status is: created, loaded-complete, open

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 Application Connectivity

Two modes, supported for all APIs (ODBC, OCI, JDBC, ODP.NET, ...)



Direct mode

- Apps run on same host as database element
- Apps directly map database shared memory (via TT engine)
- No context switches, no IPC for database access
- Ultra low latency (in process direct memory access)

Client/server mode

- TCP/IP connections between app and TT server processes
- TT Server process is a multi-threaded direct mode app
- Each interaction involves 1 or more n/w round trips
- More processes on DB host, more context switches
- More overhead, higher latency

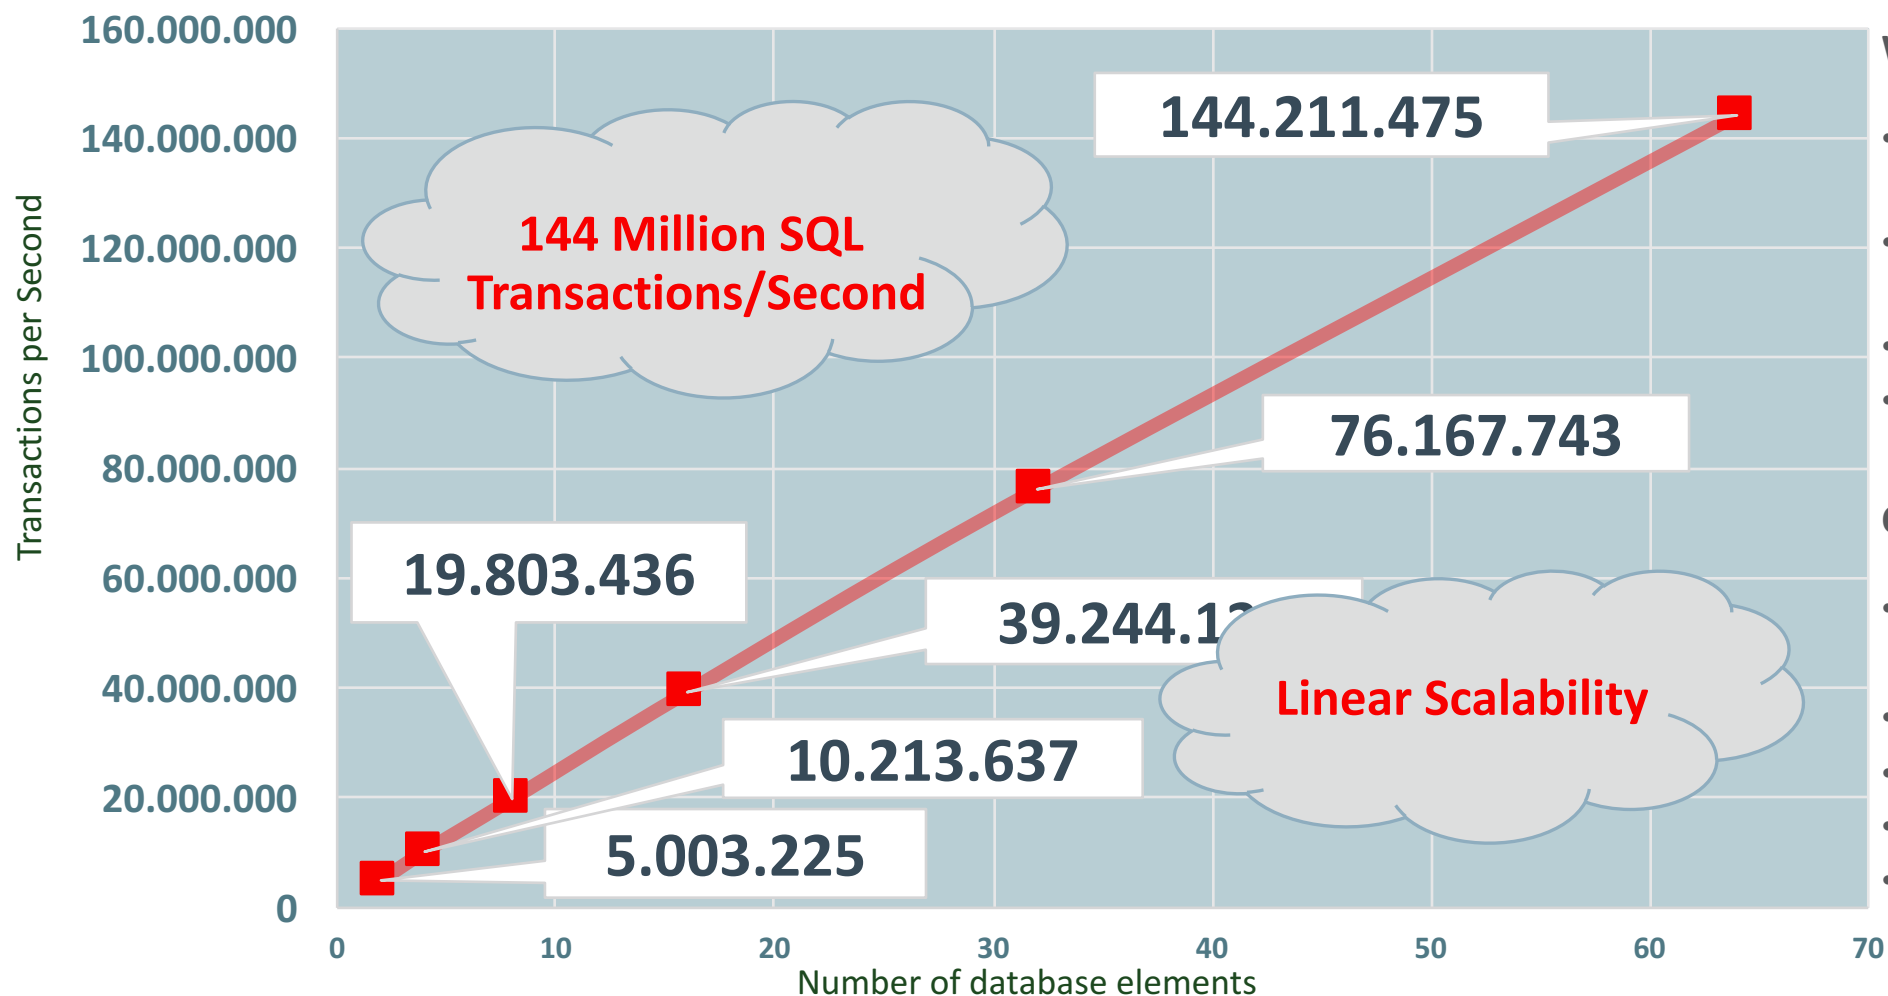
Factors influencing TimesTen Scaleout performance

For a given schema, data set and SQL workload

1.	Network	Network performance is <u>very</u> important, especially <u>latency</u>
2.	CPU, RAM, Storage	Hardware capacity and performance is important
3.	Number of machines/VMs	One (or more) Scaleout instances per machine/VM
4.	Number of Scaleout instances	One database element per instance
5.	Data distribution choices	Hash, reference, duplicate; distribution keys
6.	Local indexes	Faster data access, slower DML
7.	Global indexes	Much faster data access, much slower DML
8.	Application	Design, implementation, connectivity mode, data locality optimisations (routing API)

TimesTen Scaleout: Unmatched Transaction Throughput

SQL Transactions/sec



Workload

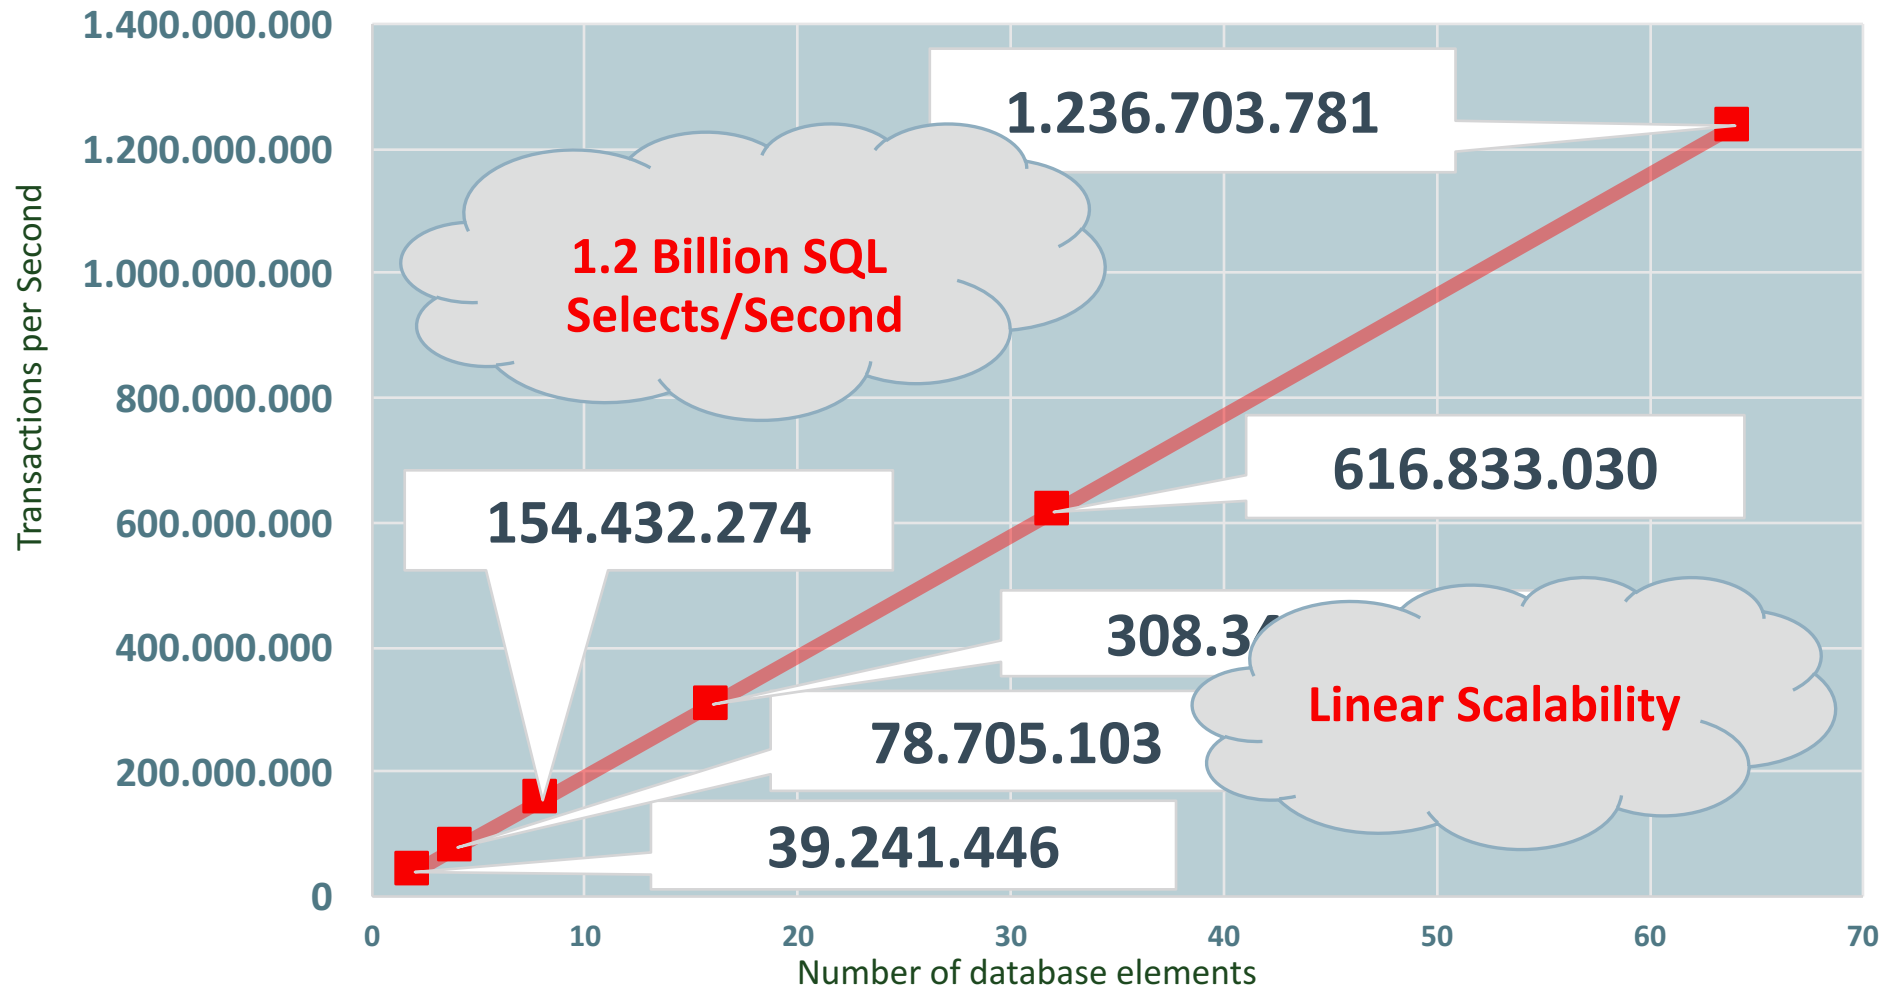
- The benchmark is an OLTP workload called **TPTBM**
- Mixed workload with **80% read / 20% update** (key based)
- **Direct mode**
- **Locality optimization**

Configuration

- Oracle Bare Metal Cloud
64 x BM.HighIO1.36
- Platform - Oracle Linux x86 64-bit
- CPU - E5-2699 v3 @ 2.30GHz
- Network - 10G Ethernet
- Storage - NVMe Disk Storage

TimesTen Scaleout: Unmatched Read Scalability

SQL Selects/sec



Workload

- The benchmark is an OLTP workload called **TPTBM**
- Workload is **100% read** (key based)
- **Direct mode**
- **Locality optimization**

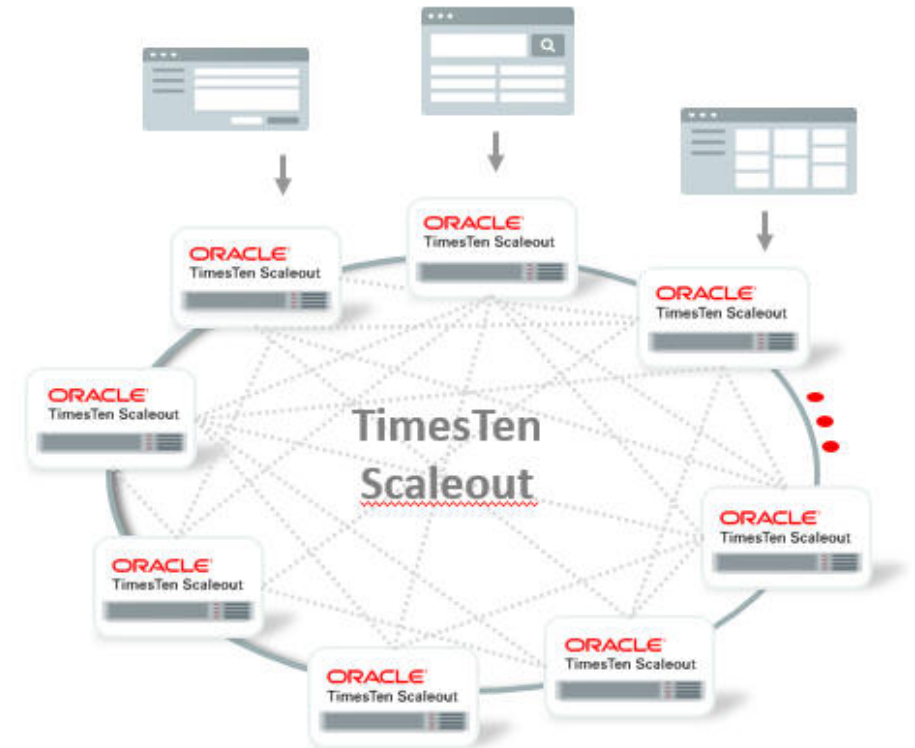
Configuration

- Oracle Bare Metal Cloud
64 x BM.HighIO1.36
- Platform - Oracle Linux x86 64-bit
- CPU - E5-2699 v3 @ 2.30GHz
- Network - 10G Ethernet
- Storage - NVMe Disk Storage

TimesTen Scaleout Summary

New feature in TimesTen 18.1 Release

- Extreme performance and scalability
- Single database image, data location transparency
- Full SQL, ACID transactions, full consistency
- 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 Velocity Scale Beta feedback

Samsung

With its elastically scalable architecture and strong HA, TimesTen Scaleout is a very attractive solution for the next generation of network systems!

-- Shin DongKeun, Principal Engineer



China Mobile

During the beta we experienced the powerful capabilities of Oracle TimesTen Scaleout with pleasure. We found it to be unique amongst similar types of products and are looking forward to it applied in our next generation distributed realtime billing system.

-- Chen Zhiheng, Business Support System Planner



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 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 and capacity


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 supports over 30 million subscribers and delivers 15 million targeted marketing messages per day.



Want to learn more?

- TimesTen OTN Portal
<http://www.oracle.com/technetwork/database/database-technologies/timesten/overview/index.html>
 - Product Information
 - Presentations, use cases, whitepapers, FAQs, ...
 - Software Downloads
 - Product Documentation
 - Scaleout Demo / Learning VirtualBox VM download
- TimesTen QuickStart and Scaleout Sample Programs on GitHub
<https://github.com/oracle/oracle-timesten-samples>
- Contact me: Chris Jenkins (chris.jenkins@oracle.com)

Q

&

A

Integrated Cloud

Applications & Platform Services

ORACLE®