

October 22–25, 2018

SAN FRANCISCO, CA

#OOW18

ORACLE
**OPEN
WORLD**

oracle.com/openworld

ORACLE®

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

Oracle TimesTen Scaleout Getting Started

ORACLE
OPEN
WORLD

Tirthankar Lahiri
Vice President, Data & In-Memory Technologies
Oracle Development

Chris Jenkins
Senior Director, Product Management
Oracle Development

October 23, 2018

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, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Database Platform of the **Future** – *Complete and Integrated*

Services
Any Application
Any Data

- OLTP, IoT 
- Real Time Analytics 
- Big Data 
- Web, Micro Services 
- Data Science, AI, ML 
- Spatial, Graph, Text, Media 



Architecture
Fast, Low Cost, Secure
Scalable, Available

- In-Memory Database** 
- In-Database Multitenancy 
- In-Database Security 
- Fault-Tolerant Scale-Out 

Systems
Engineered
for Databases

- 
- Optimized Compute 
- Smart Storage 
- Fastest Networking 


Cloud
Autonomous, Elastic
Wherever You Want

- Public Cloud 
- Cloud At Customer 

In-Memory Database Technology **Facts**

The Time is NOW

- Next Generation Enterprises **must** be real-time
 - In-Memory is **essential** for real-time processing
 - In-Memory is **essential** in next-gen database platforms
- Current Memory sizes enable In-Memory **today**
 - Oracle X7-2 server: **Up to 1.5TB** of DRAM
 - Not using in-memory is like using Windows 98 in 2018



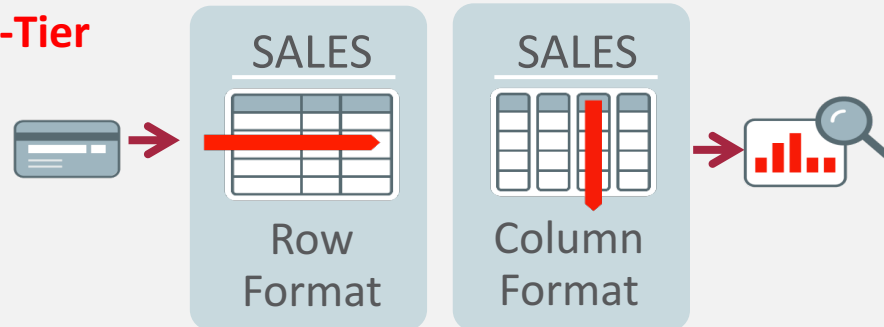
In-Memory Across Tiers in the Oracle Database Stack

Application-Tier



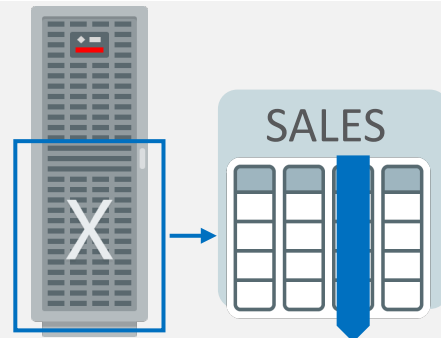
- **TimesTen In-Memory Database**
 - Latency Critical custom OLTP applications
 - **Microsecond** response time
 - Standalone or Cache for Oracle Database

Database-Tier



- **Database In-Memory**
 - Dual Format In-Memory Database
 - **Billions of Rows/sec** analytic data processing
 - **2-3x** Faster Mixed Workloads

Storage-Tier



- **In-Memory on Exadata Storage**
 - In-memory column format on Exadata Flash Cache
 - **5-10x** faster smart scan in storage
 - **15x** increase in total columnar capacity


Oracle TimesTen In-Memory Database

Relational Database



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

Persistent and Recoverable




- Database and Transaction logs persisted on local disk or flash storage

Extremely Fast



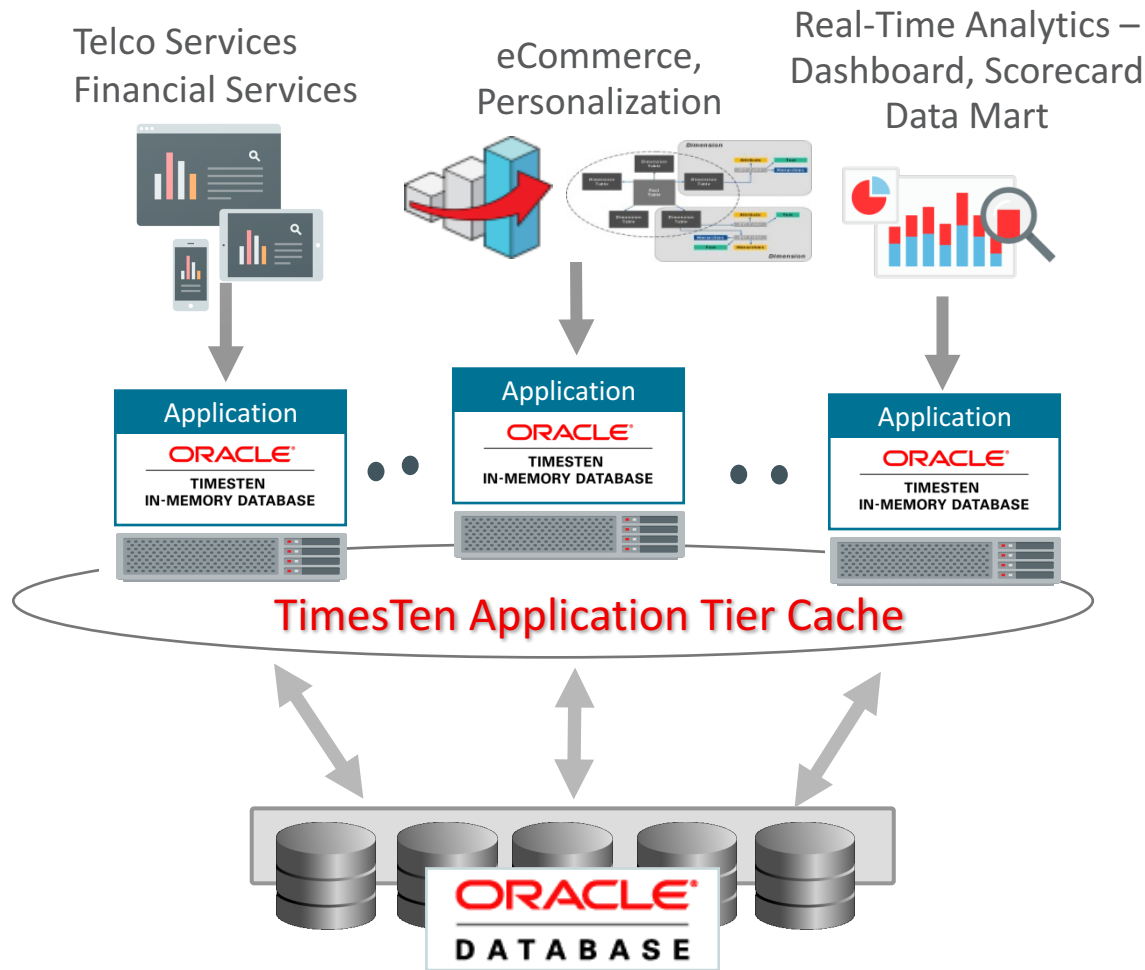
- Microseconds response time
- Very high throughput

Highly Available



- Active-Standby and multi-master replication
- Very high performance parallel replication

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
- HA and fault tolerance in the application-tier




Most Widely Used Relational In-Memory Database

Deployed by thousands of Companies



New in 18.1 Release: TimesTen Scaleout In-Memory Database

Proven TimesTen technology with Scale-Out for High Performance, and Multi-Copy for High Availability



- For High-Velocity **Extreme OLTP** applications
 - IoT, trading, telecommunications, click stream, billing, orders, fraud detection, etc.
- Performance-Oriented Design
 - Pure In-Memory, Full SQL, Full ACID Transactions
 - Fault-Tolerant Scale-Out
 - Sophisticated and Parallel SQL for reporting and batch
- Extremely easy to install and operate
 - **Less than 15 mins** to install, deploy, and run application

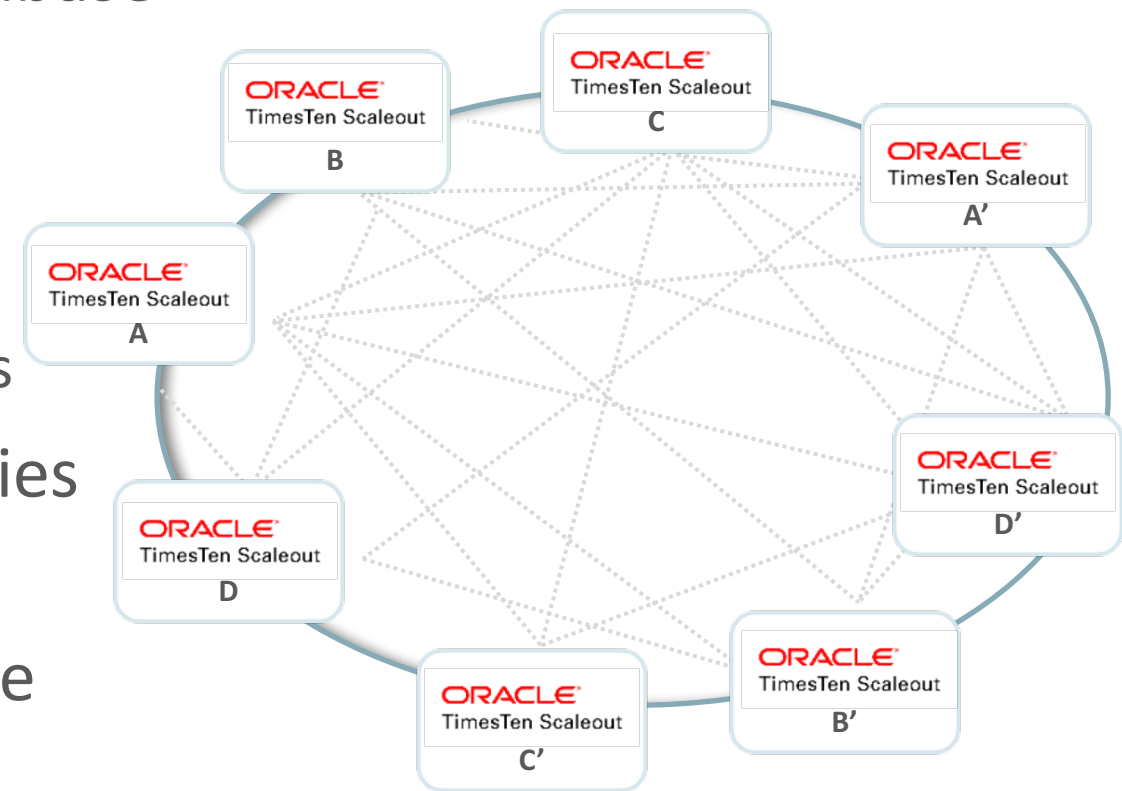
Getting Started with TimesTen Scaleout

TimesTen Scaleout: The World's Fastest OLTP Database

Distributed, Shared Nothing, In-Memory Database

Single-Image Database with High Availability and Elasticity

- Appears to applications as a single database
 - **Not** as a sharded database
- Scale-out and scale-in
 - Data automatically redistributed
 - Workload automatically uses new elements
- Built-in HA via multiple fully-active copies
 - Copies automatically kept in sync
- Highly compatible with Oracle Database
 - Data types, APIs, SQL & PL/SQL



Distribution Methods for All Scenarios

Distribute Table Data by **Hash**, **Reference** or **Duplicate**

- Distribute by **Hash**

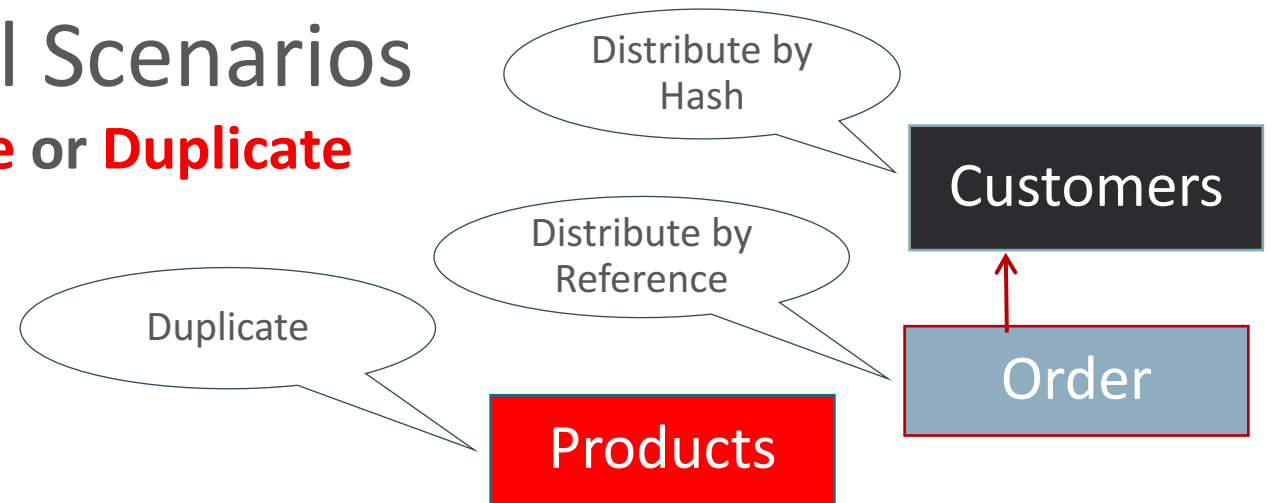
- Primary key or user-specified columns
- Consistent hash algorithm
- Examples: Customers, Subscribers, Accounts

- Distribute by **Reference**

- Co-locate related data to optimize joins
- Based on FK relationship
- Supports multi-level hierarchy

- Full **Duplicate**

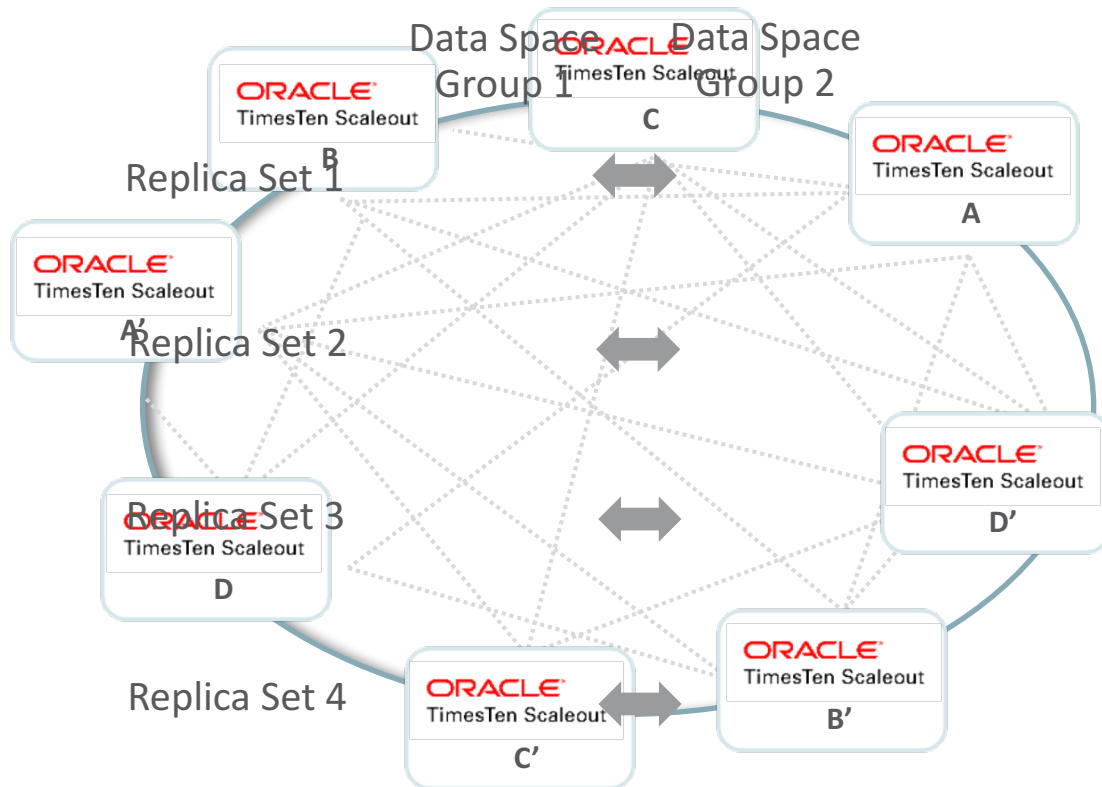
- Identical copies on all elements
- Useful for reference tables
- Read and join optimization



Element 1		Element 2		Element 3		Element 4			
0	David	1	Bill	2	Olaf	3	Chi		
4	Igor	5	Sam	6	Henri	7	Simon		
8	Tim	9	Charles	10	Jie	11	Chris		
1	0	16/6/15	2	5	16/2/22	3	3	16/3/1	
6	8	16/3/22		5	6	16/5/10	4	11	16/2/5
phone	100	phone	100	phone	100	phone	100		
tablet	200	tablet	200	tablet	200	tablet	200		
watch	300	watch	300	watch	300	watch	300		

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

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 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



TimesTen Grid

Status Database Definition Topology

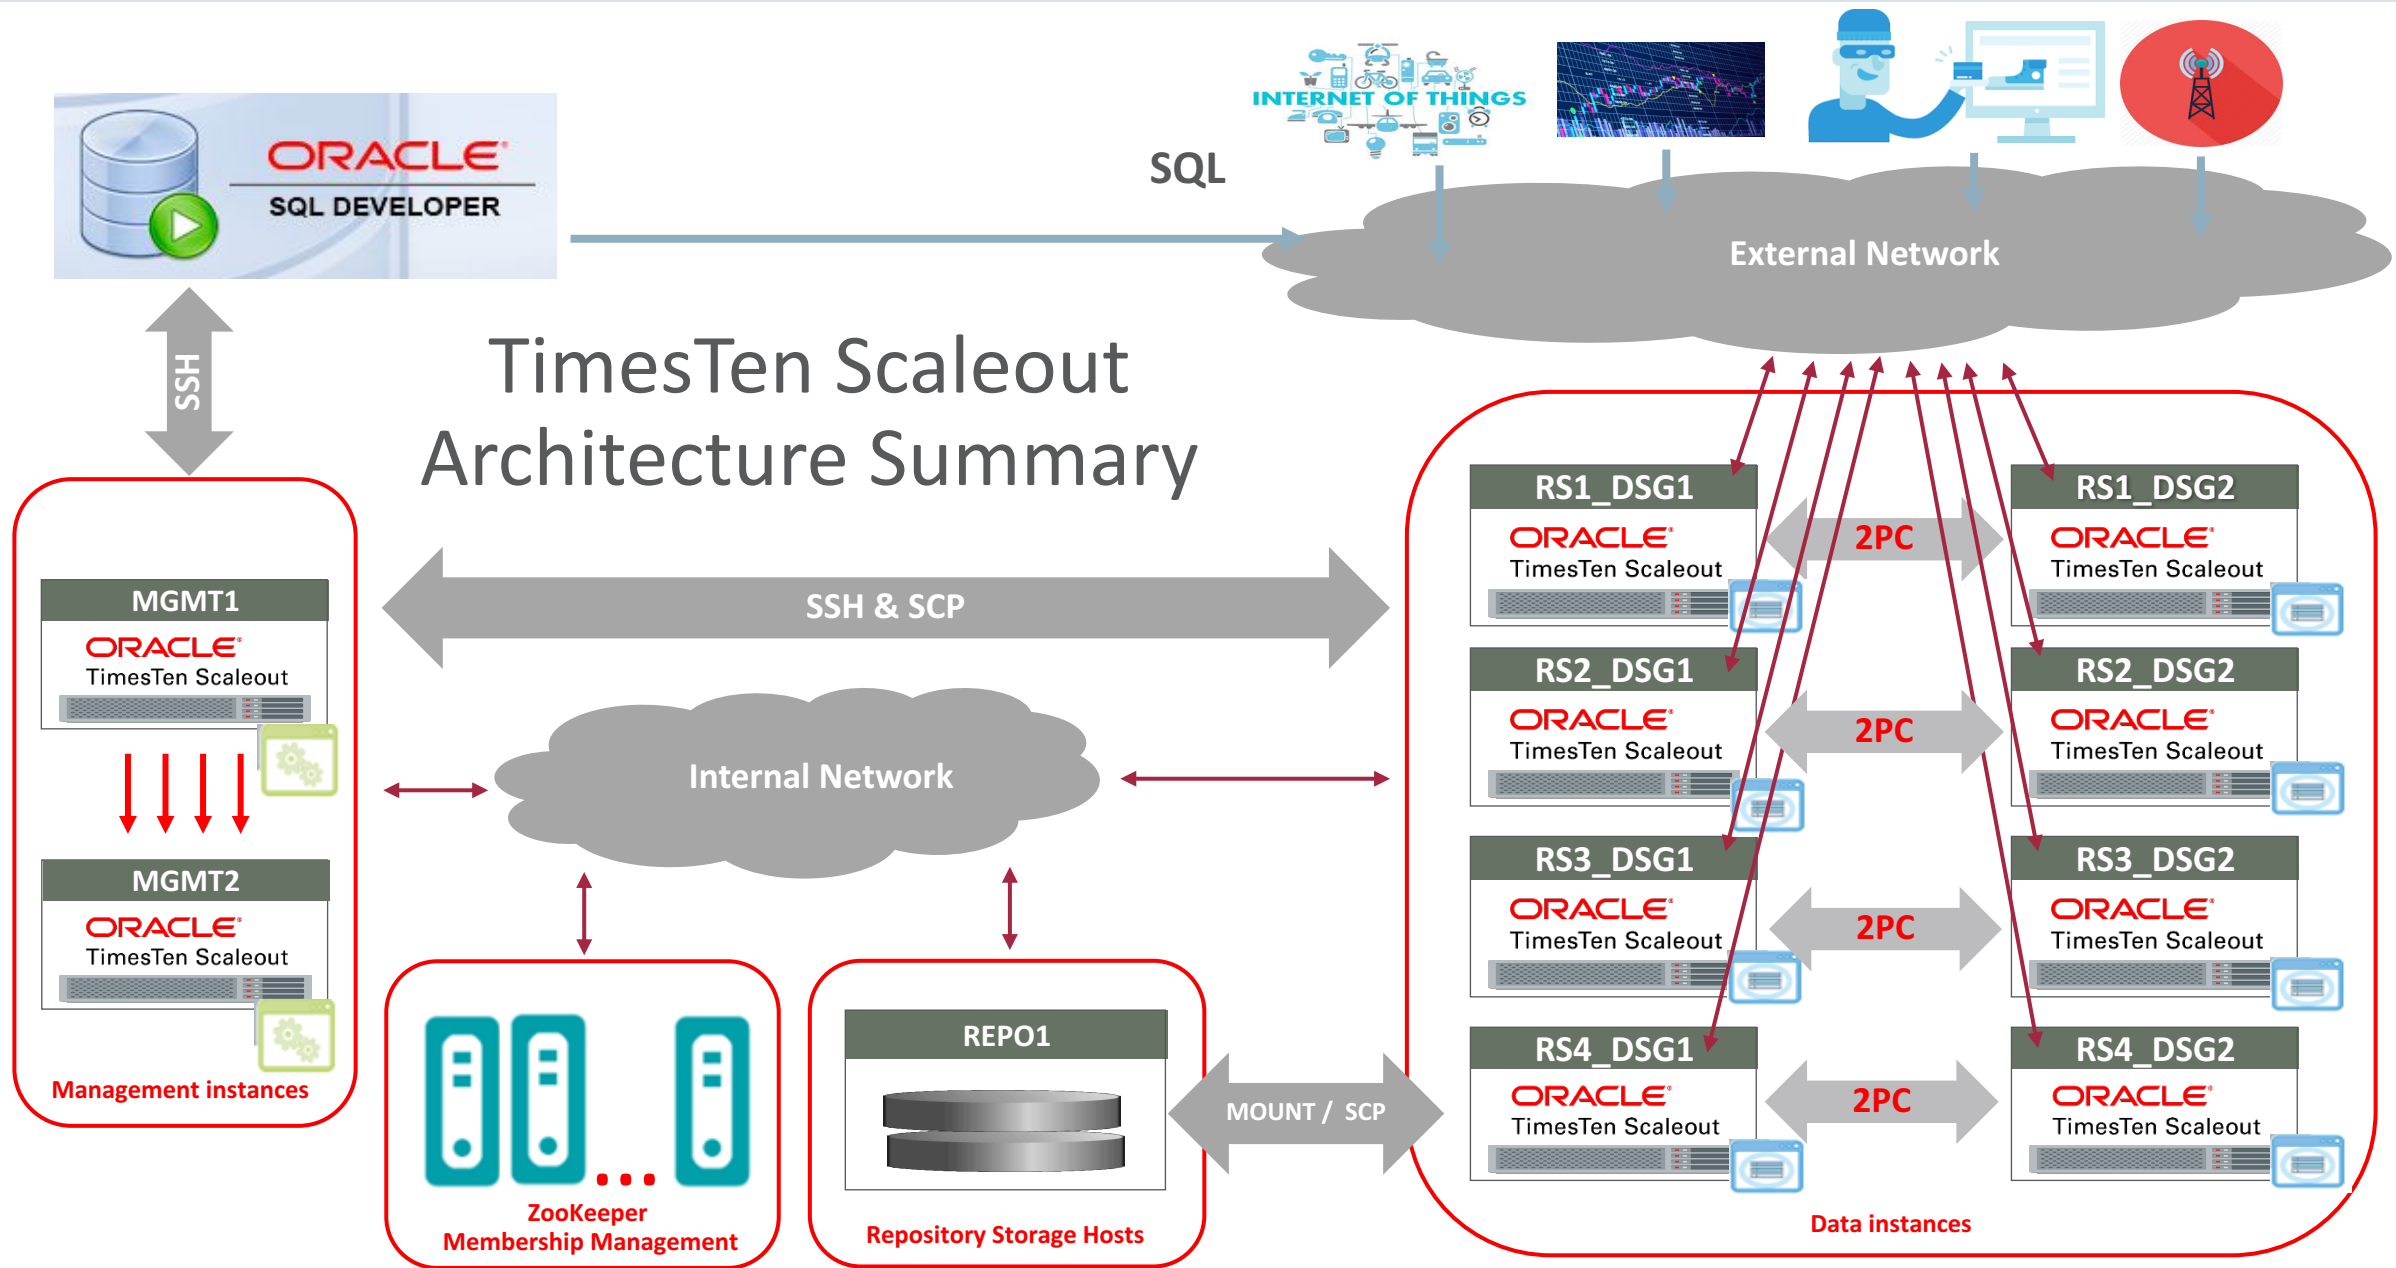
Database demodb status is: created, loaded-complete, open

Number of application connections to demodb: 0

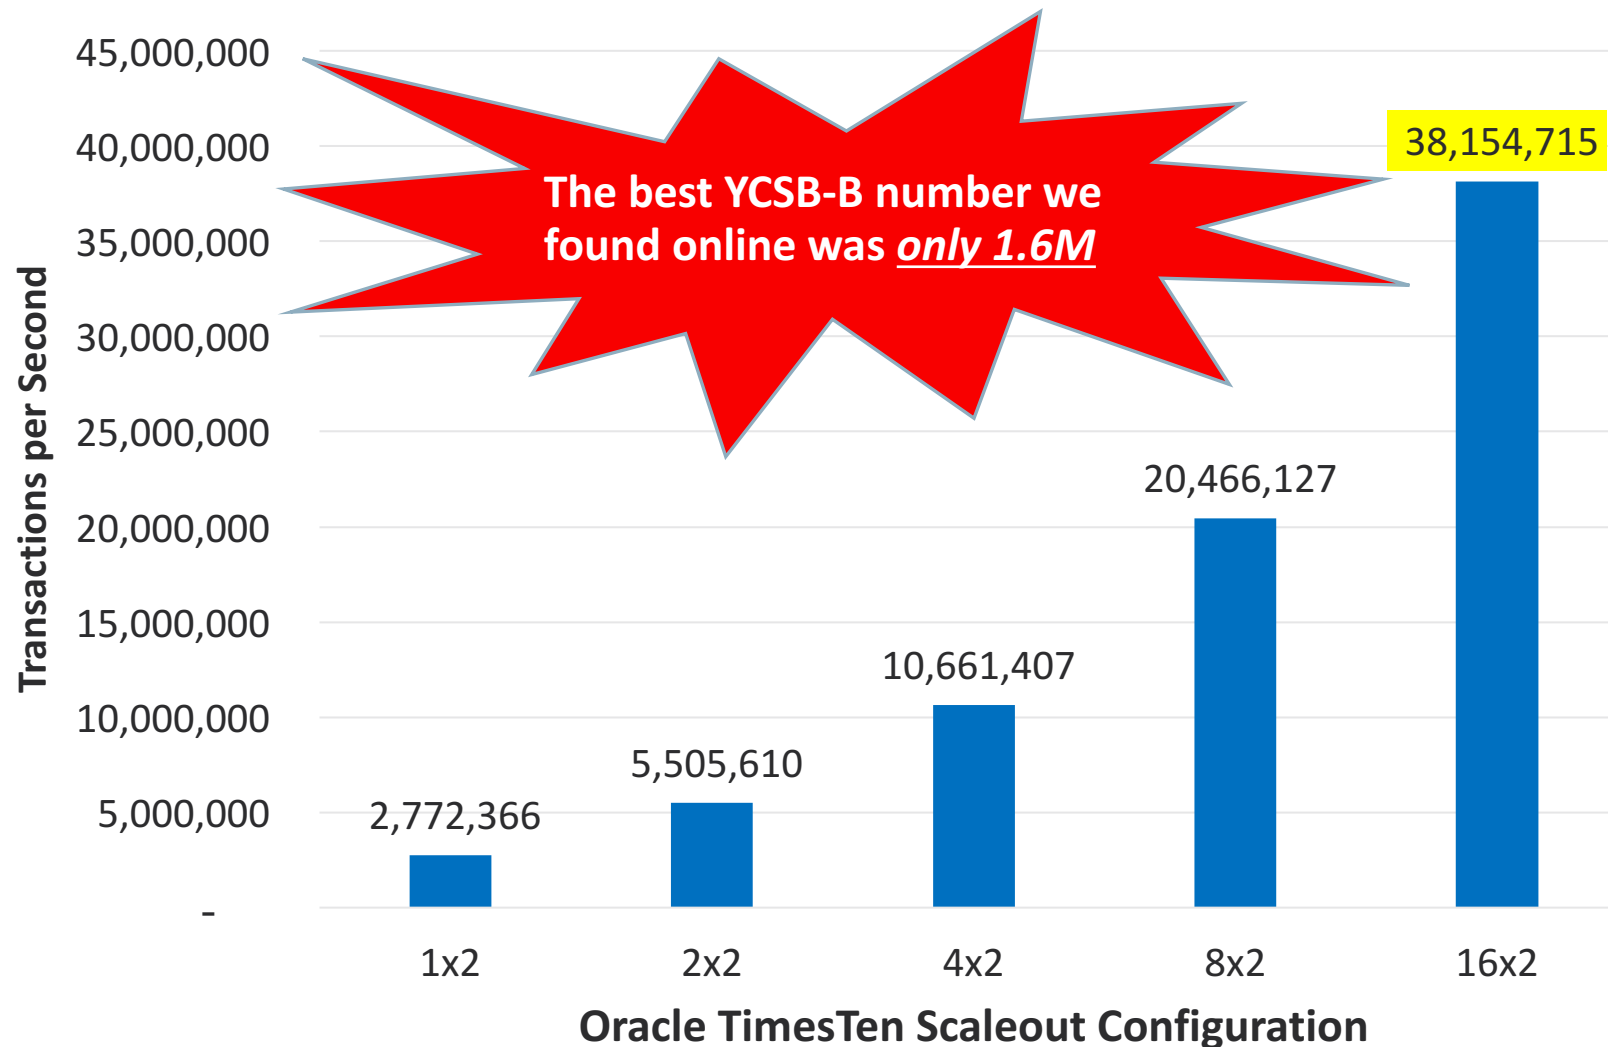
Number of system connections to demodb: 168

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



YCSB Workload B (95% Read 5% Update): **38 Million TPS**



YCSB version 0.15.0

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

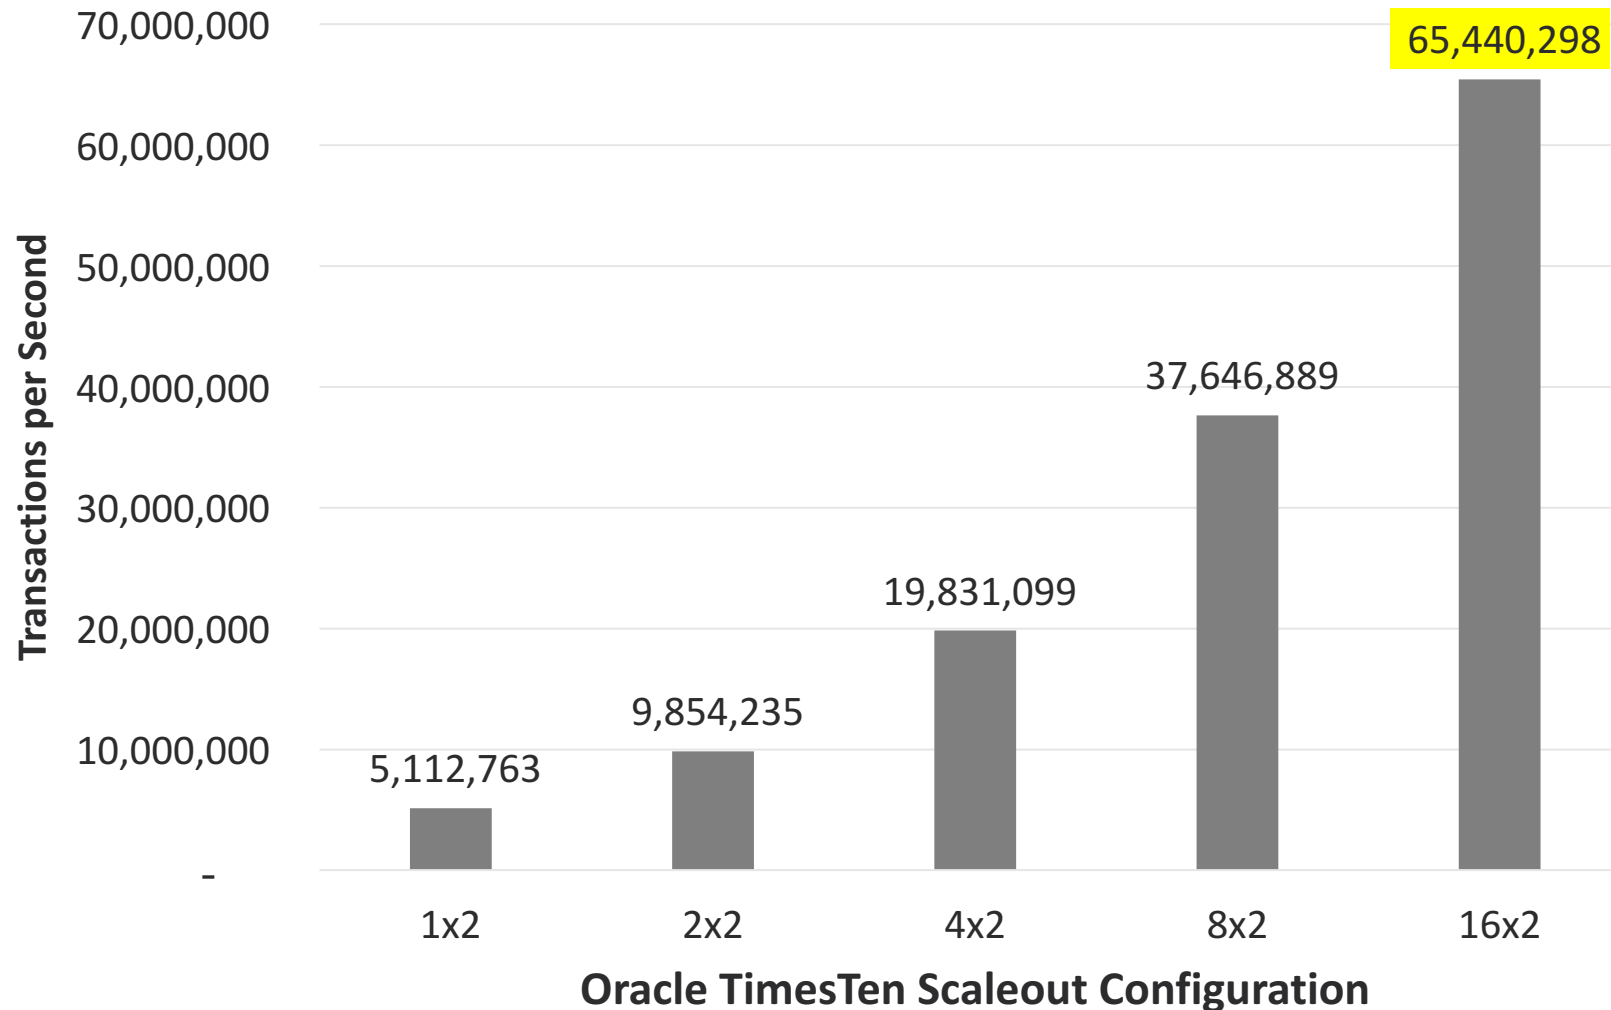
TimesTen Scaleout

- 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 TPS**



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

TPTBM 80% Read 20% Update: **153 Million TPS**



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

Planning: Example Hardware Configurations

Dev/Test versus Production

Dev/test grid (minimal)

- 1 management host
 - ZooKeeper server
 - Management instance
 - Repository (optional)
- 2 data hosts
 - 2x1 grid – No HA
 - 1x2 grid – HA
- 1 GbE dedicated network
 - 10 GbE is much better

Production grid

- 2 management hosts
 - ZooKeeper servers
 - Management instances
- Repository host
 - ZooKeeper server
- N data hosts: N depends on
 - Required storage capacity
 - Required processing capacity
 - HA or no HA (HA strongly recommended)
- Dedicated 10 GbE (or better) internal network

Optimizing: Performance

Internal Network

- Minimize latency, maximize bandwidth
- 10 GbE is the minimum
- 25 GbE+ is better

External Network

- Depends on
 - Application requirements
 - SLAs
 - ...
- In general faster is better

Hosts

- Physical hosts better than VMs
- Fewer large hosts better than many small hosts

Data model and schema

- Table distribution choices
 - Distribution type
 - Distribution key
- Indexing
 - Local versus global

Optimizing: Availability

K-factor

- K=1 – no HA!
- K > 1 recommended for production

Data Space Groups

- A DSG = a set of hosts that might all fail at the same time
 - Rack with common PSU
 - Cloud AD / AZ
 - ...
- Assignment of hosts to DSGs is critical for maximizing availability

Tradeoffs

- Increasing the K-factor
 - Increases availability
 - Reduces performance for DML
- Separating DSGs (e.g. Cloud AD / AZ)
 - Increases availability
 - Increases latency between DSGs
 - Reduces performance for DML
- As always it is a balance based on your requirements

Getting Started: **On Premises**

- Prepare hardware/VMs and network
 - Prepare OS, admin user and filesystem locations on all hosts
 - Download TimesTen to management host
 - Unzip download to create an installation
- Prepare passwordless SSH
 - Use **ttGridAdmin gridSSHConfig**
 - Prepare configuration files for ttGridRollout
 - Deploy your TimesTen Scaleout Grid and Database
 - One command; **ttGridRollout**
 - Have fun with your on-premises Scaleout environment!

Getting Started: Oracle Cloud (OCI) using BYOL

- Prepare your OCI tenancy
 - Account
 - Resources
 - Prepare your 'client' system
 - Linux x86 or macOS
 - Install Terraform
 - Download TimesTen and Java packages to client system
 - Download whitepaper (OTN) and deployment scripts (GitHub)
- Extract scripts on 'client'
 - Customize basic configuration
 - See the whitepaper
 - Use scripts to configure OCI systems and deploy TimesTen Scaleout
 - Just 4 commands!
 - Enjoy your Cloud Scale IMDB!

Getting Started: Other resources

- TimesTen Blog

<https://blogs.oracle.com/timesten>

TimesTen Portal

<https://www.oracle.com/database/technologies/related/timesten.html>

- Documentation

<https://docs.oracle.com/database/timesten-18.1/>

- QuickStart and Samples

<https://github.com/oracle/oracle-timesten-samples>

- TimesTen Scaleout VirtualBox VM

<https://www.oracle.com/technetwork/database/database-technologies/timesten/downloads/timesten-181-vm-download-4480199.html>

TimesTen Scaleout HA Demo

TimesTen Scaleout: The World's Fastest OLTP Database

Demo: High Availability in TimesTen Scaleout

Scaleout grid

- 12 data hosts (Cloud VMs)

Database

- K=2, 6 replica sets (6x2)

Workload

- 80% read, 20% write

Application

- Running on a separate host
- Client/server connectivity

Demo sequence

- Run workload
- Kill one of the data instances
- Recover and resync the failed instance
- See what happens when we do the above!

Wrap-up

TimesTen Scaleout: The World's Fastest OLTP Database

TimesTen Scaleout: Scalability Without Compromise

- Traditionally, relational databases were considered sophisticated but hard to scale
 - As a result many users adopted NoSQL / Key-Value solutions to achieve scalability
 - But gave up most of the important functionality of relational databases: SQL, transactions, integrity constraints, consistency, ad-hoc analytics, etc.
- TimesTen Scale-Out offers a scale-out solution without these compromises
 - Built on mature, proven TimesTen In-Memory Database as a foundation
 - Extreme performance while providing SQL, full ACID transactions, Indexes, etc.
 - Built in fault tolerance for cloud-scale workloads
 - Easy to deploy and manage, even for large grids

Other TimesTen Sessions

Code	Title	When	Where
PRO4017	Oracle TimesTen Scaleout: Developing Applications for OLTP and IoT	Wed Oct 24 @ 11:15 am	Moscone West 3001
HOL6319	Next-Generation OLTP: Oracle TimesTen Scaleout	Tue Oct 23 @ 5:15 pm Wed Oct 24 @ 12:45 pm	Marriott Marquis (Yerba Buena Level) - Salon 3/4
ESS6915	Next-Gen Cloud Technologies That Accelerate Business Transformation (Intel Keynote)	Tue Oct 23 @ 12:30 pm	Yerba Buena Center for the Arts (YBCA) Theater
Demo Grounds	Booth DBA-A04		

Q

&

A